E-crime

CS642: Computer Security

Professor Ristenpart

http://www.cs.wisc.edu/~rist/

rist at cs dot wisc dot edu

ATTN; Thomas Ristenpart

ALan Daniel <alandaniel33@yahoo.co.uk> to crypto2015@iacr.org

Confidant: Gregor Leander

I am Barr. Allan Daniel (retired).

I have a knowledge of an unclaimed estate proceed left in a codicil by a deceased Veteran, Scott G. Ristenpart....{2nd Battalion,505th Parachute infantry Regiment,3rd Brigade combat Team,82nd Airborne Division at Ft Bragg NC..}.

I am contacting you to seek for your co-operation to stand in as a possible heir since you have the same last name as the deceased beneficiary so that the proceed of this estate could be taken as a fortune before the Probate ref:no is revoked for a lack of "Executor/Executrix".

If you agree to this confidential consultation....I will advise on the programmed of work to expedite this claim through a legitimate arrangement for a mutual benefit ..

Thank you for your anticipated response.

Best regards,

Allan Daniel Jr.

Spam, phishing, scams

- Spam
  - unsolicited bulk emails
  - 2006: 80% of emails on web, 85 billion messages a day
  - 2009: 95% of emails blocked as spam (100s of billions)
- Scam spam
  - Nigerian emails (advanced fee fraud / confidence trick)
- Phishing
  - trick users into downloading malware, submitting CC info to attacker, etc.
  - Spear phishing: targeted on individuals (used in high-profile intrusions)

Spanish Prisoner confidence trick

- 19th century
- In contact with rich guy in Spanish prison
- Just need a little money to bribe guards, he'll reward you greatly

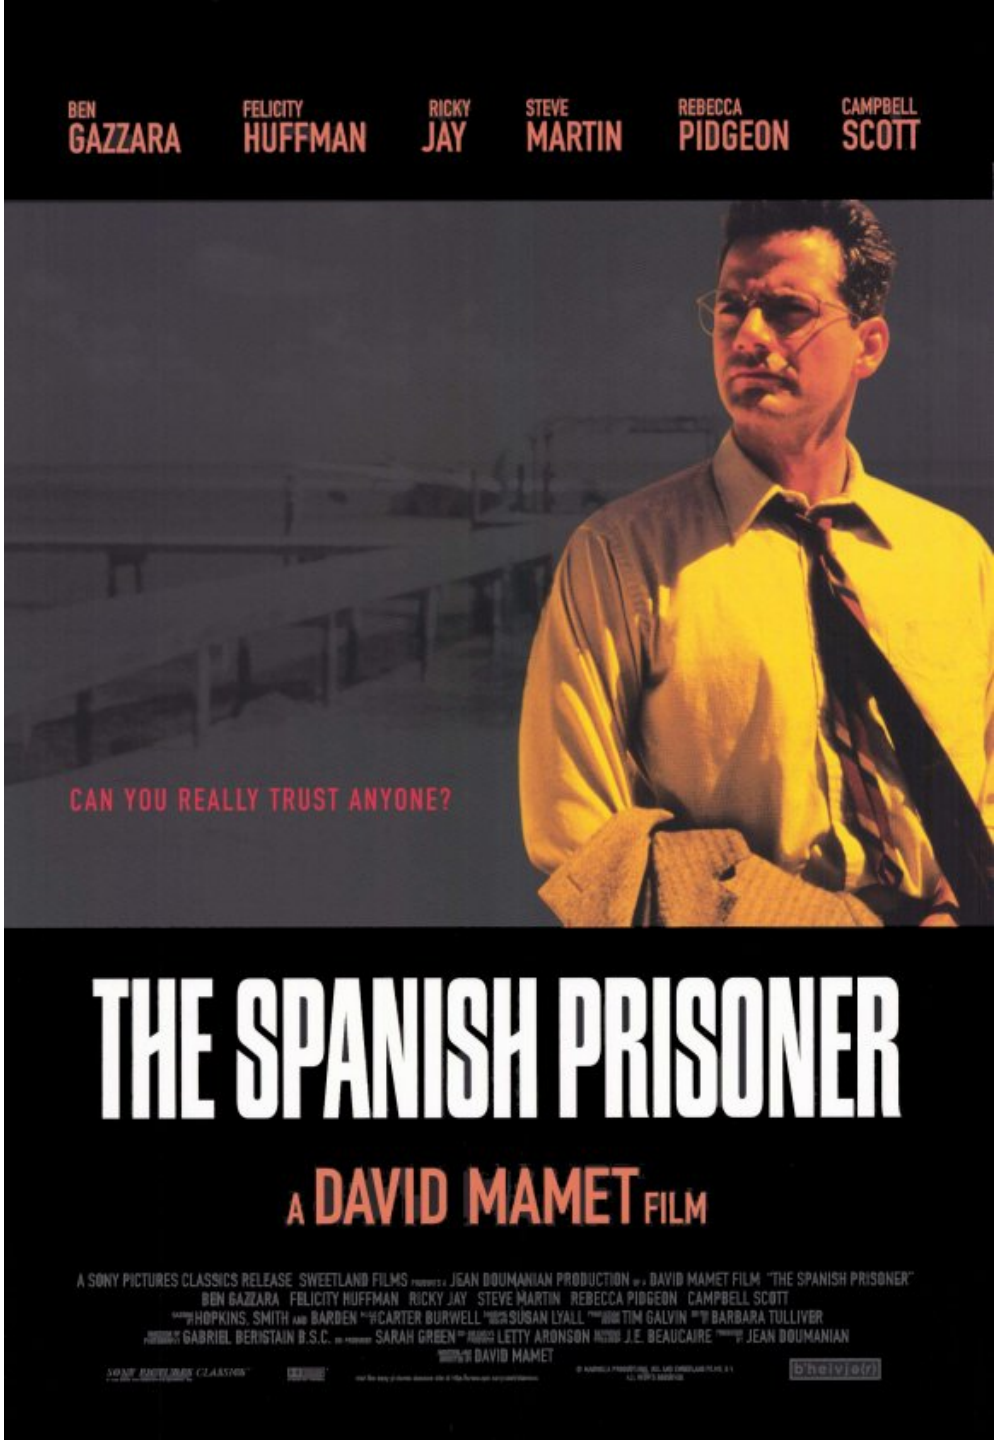

from Mrs. Zarina Al-Usman <zarina_alusman@kimo.com>

subject **Re: My Desire for you Over Less Privileged Children**

to undisclosed-recipients: ;

Hi Dear,

I am Mrs. Zarina Al-Usman, I have been diagnosed with Esophageal cancer .It has defied all forms of medical treatment, and Right now, I have only about a few months to live and I want you to Distribute my funds worth Twelve Million Five Hundred Thousand US Dollars to charities homes in your country.

I have set aside 40% for you and your family so keep this as a secret to yourself because this will be my last wish.
Yours Truly,

Mrs. Zarina Al-Usman

---

from Fatemeh Akhoundi <akhounf@st
subject **Your mailbox has exceeded its**
to undisclosed-recipients:;

Junk Mail

Your mailbox has exceeded its
99.7% of its Quota limit of 10
you have updated your webmail
copy the link below and paste

http://upgradeportal11.media.o

We are sincerely sorry for any
to serve you better.
Thanks for your co-operation.
Webmail Update Team

csehelp@cs.ucsd.edu to tristenp@cs.ucsd.edu

This is an automatic notification of your current disk space usage on the CSE mail server:

csemailbox.ucsd.edu

Your account status:

Current utilization: 95.33%

Space used:	976 MB
Available space:	47 MB
Account limit:	1024 MB

Once your quota has been reached, mail will no longer be delivered to your account, and will be returned to the sender as undeliverable.

If you are not sure where to look for mail that can likely be deleted to clear space in your account, you may likely have large amounts of mail in your Trash and/or Junk folders. Also, you may have a large amount of mail accumulating in your Sent folder over time, if you have configured your mail client to automatically save sent messages.

Your account limit may be increased for an additional charge, as per the CSE Recharge Policy. Please contact CSEHelp regarding quota increases.

Please reply to this message or contact CSEHelp <csehelp@cs.ucsd.edu> if you have any questions or require assistance.

Thank you,

Spam volume

Billions of messages

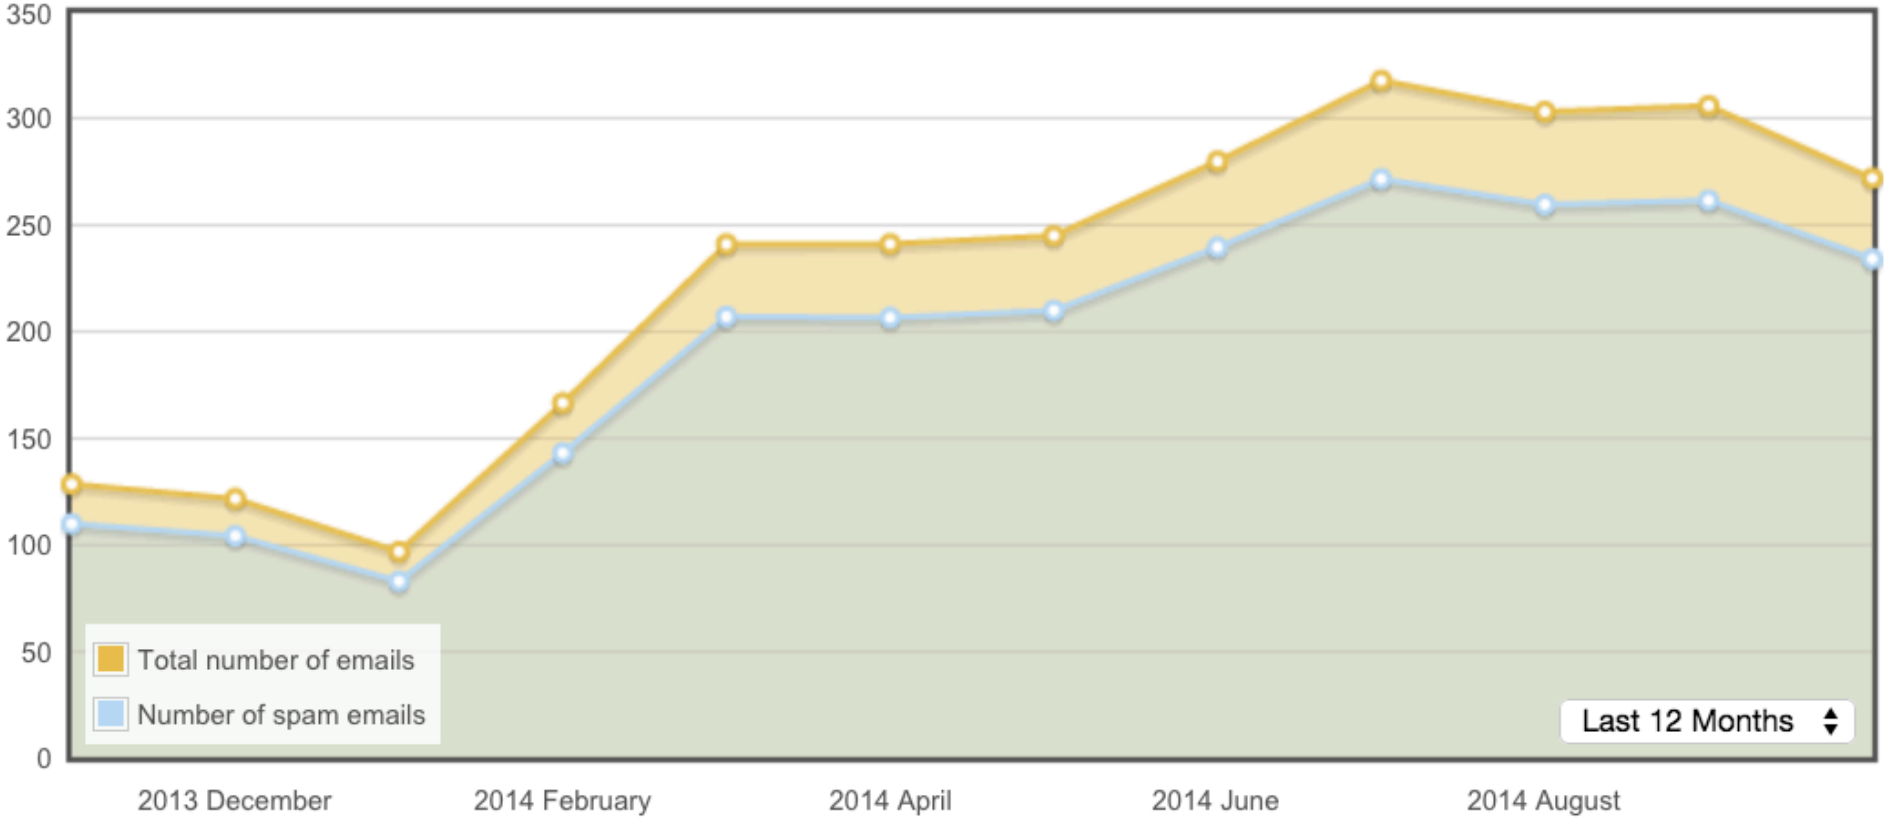

http://www.senderbase.org/static/spam/#tab=1

Spam

- The frontend (email recipients)
  - Filtering, classification
  - Psychology
- The backend (email generation)
  - Open email relays
  - Botnets
  - Social structure
    - Affiliates
    - Criminal organizations

Spam Classifiers

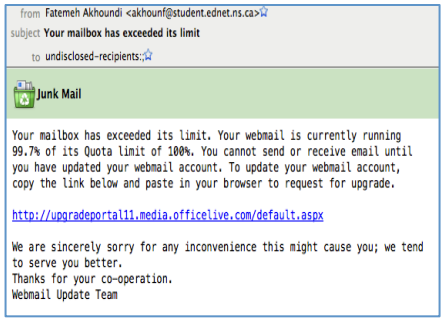

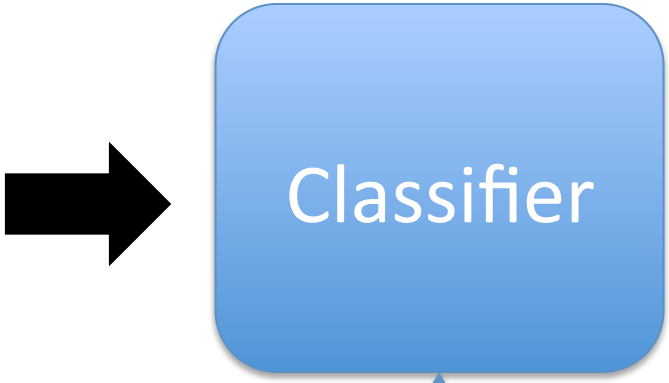

Spam / Ham

This classifier will be trained from a large corpus of labeled data

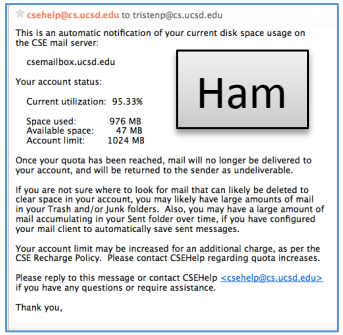

Ham

Spam

...

Naïve Bayes Classifier

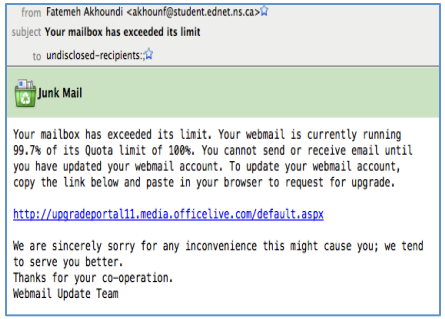
from Fatemeh Akhoundi <akhounf@student.ednet.ns.ca>
subject Your mailbox has exceeded its limit
to undisclosed-recipients:;

Junk Mail

Your mailbox has exceeded its limit. Your webmail is currently running 99.7% of its Quota limit of 100%. You cannot send or receive email until you have updated your webmail account. To update your webmail account, copy the link below and paste in your browser to request for upgrade.

http://upgradeportal11.media.officelive.com/default.aspx

We are sincerely sorry for any inconvenience this might cause you; we tend to serve you better.
Thanks for your co-operation.
Webmail Update Team

Represent email as "bag of words"

quota	1	x_1
webmail	4	x_2
wisconsin	0	x_3
viagra	0	x_4
⋮	⋮	

Intuition: spam and ham have different distribution of keywords

$$\Pr[\text{spam} \mid x_1, x_2, \ldots, x_n] = \frac{\Pr[x_1, x_2, \ldots, x_n \mid \text{spam}]\ \Pr[\text{spam}]}{\Pr[x_1, x_2, \ldots, x_n]}$$

Bayes' theorem

$$= \frac{\Pr[\text{spam}] \prod \Pr[x_i \mid \text{spam}]}{\Pr[x_1, x_2, \ldots, x_n]}$$

"Naïve": assume words independent

$$\Pr[\text{ham} \mid x_1, x_2, \ldots, x_n] = \frac{\Pr[\text{ham}] \prod \Pr[x_i \mid \text{ham}]}{\Pr[x_1, x_2, \ldots, x_n]}$$

Naïve Bayes Classifier

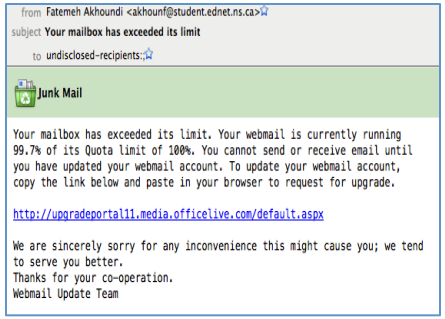

Represent email as "bag of words"

quota	1	x_1
webmail	4	x_2
wisconsin	0	x_3
viagra	0	x_4
⋮	⋮	

Intuition: spam and ham have different distribution of keywords

$$\Pr[\text{ spam } | \; x_1, x_2, \ldots, x_n] = \frac{\Pr[\text{spam}] \prod \Pr[x_i \mid \text{spam}]}{\Pr[x_1, x_2, \ldots, x_n]}$$

$$\Pr[\text{ ham } | \; x_1, x_2, \ldots, x_n] = \frac{\Pr[\text{ham}] \prod \Pr[x_i \mid \text{ham}]}{\Pr[x_1, x_2, \ldots, x_n]}$$

Classify as spam if: $\Pr[\text{ spam } | \; x_1, x_2, \ldots, x_n] > \Pr[\text{ ham } | \; x_1, x_2, \ldots, x_n]$

Naïve Bayes Classifier

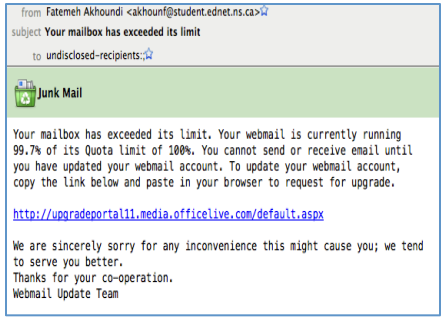

Represent email as "bag of words"

quota	1	x_1
webmail	4	x_2
wisconsin	0	x_3
viagra	0	x_4
⋮	⋮	

Intuition: spam and ham have different distribution of keywords

Estimate these from labeled training data

$$\Pr[\text{spam} \mid x_1, x_2, \ldots, x_n] = \frac{\Pr[\text{spam}] \prod \Pr[x_i \mid \text{spam}]}{\Pr[x_1, x_2, \ldots, x_n]}$$

$$\Pr[\text{ham} \mid x_1, x_2, \ldots, x_n] = \frac{\Pr[\text{ham}] \prod \Pr[x_i \mid \text{ham}]}{\Pr[x_1, x_2, \ldots, x_n]}$$

Classify as spam if: $\Pr[\text{spam}] \prod \Pr[x_i \mid \text{spam}] > \Pr[\text{ham}] \prod \Pr[x_i \mid \text{ham}]$

Spam classifiers

- Real classifiers more complex than this
  - Other features: Who is sender? How many links embedded? Is it from an open mail relay?
  - Can update in real-time given labelings by user
  - For larger orgs, can leverage wide view across many email recipients
- Nowadays some companies do pretty good job of making sure spam doesn't hit your inbox
  - 95% of email gets filtered as spam (2009, ENISA Spam Survey)

Spam

- The frontend (email recipients)
  - Filtering, classification
  - Psychology, usability
- The backend (email generation)
  - Open email relays
  - Botnets
  - Social structure
    - Affiliates
    - Criminal organizations

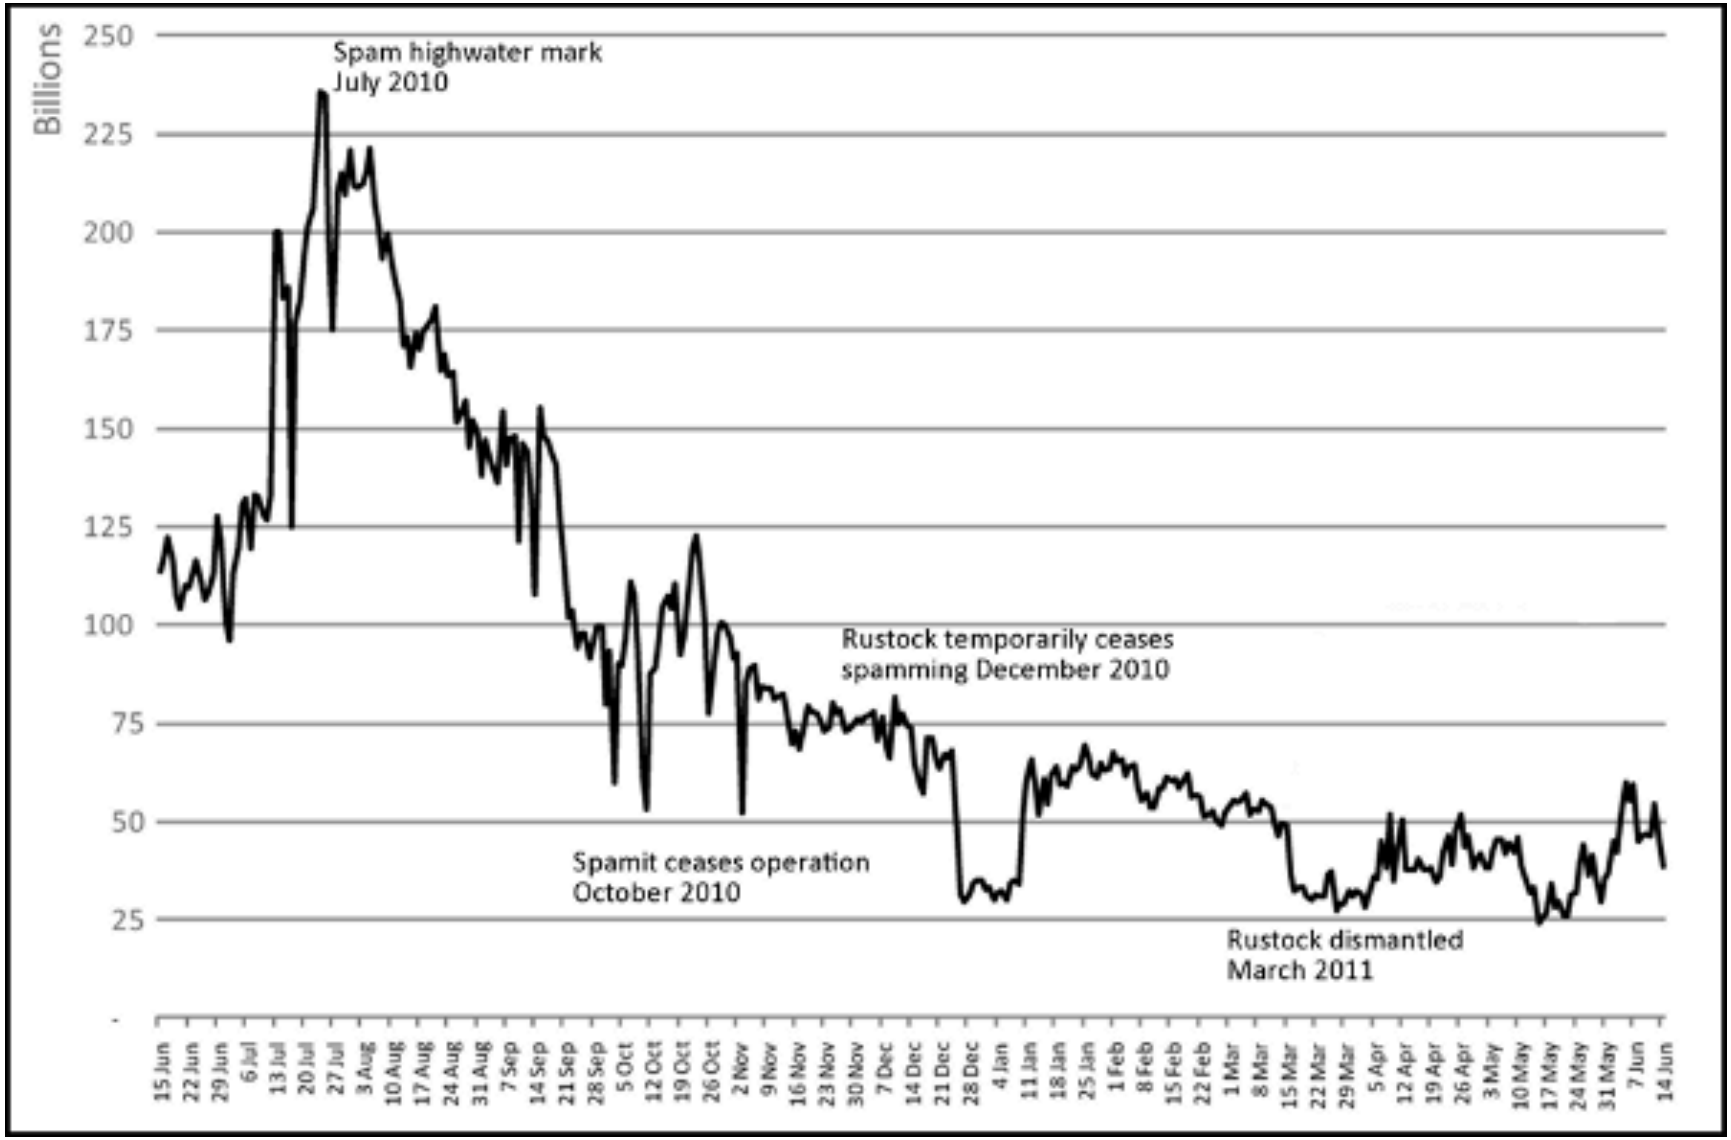

http://www.symantec.com/connect/blogs/why-my-email-went

Botnets

- Botnets:
  - Command and Control (C&C)
  - Zombie hosts (bots)
- C&C type:
  - centralized, peer-to-peer
- Infection vector:
  - spam, random/targeted scanning
- Usage:
  - What they do: spam, DDoS, SEO, traffic generation, …

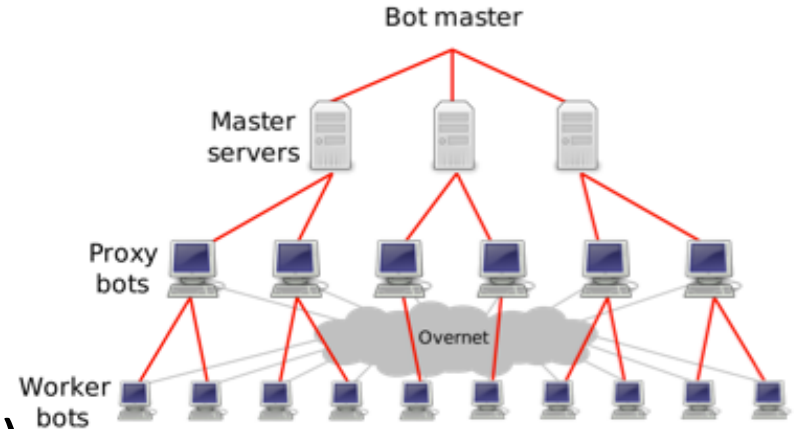

Figure 1: The Storm botnet hierarchy.

How to make money off a botnet?

- Rental
  - "Pay me money, and I'll let you use my botnet… no questions asked"
- DDoS extortion
  - "Pay me or I take your legitimate business off web"
- Bulk traffic selling
  - "Pay me to direct bots to websites to boost visit counts"
- Click fraud, SEO
  - "Simulate clicks on advertised links to generate revenue"
  - Cloaking, link farms, etc.
- Theft of monetizable data (eg., financial accounts)
- Data ransom
  - "I've encrypted your harddrive, now pay me money to unencrypt it"
- Advertise products

Underground forums

Category	Threads B	Threads S	Users B	Users S	Top Subcategory
payments	8,507	8,092	1,539	1,409	paysafecard
game-related	2,379	2,584	924	987	steam
accounts	2,119	2,067	850	974	rapidshare
credit cards	996	1160	467	566	unspecified cc
software/keys	729	1410	422	740	key/serial
fraud tools	652	1155	363	601	socks
tutorials/guides	950	537	562	393	tutorials
mail/drop srvs	751	681	407	364	packstation
merchandise	493	721	264	404	ipod
services	266	916	176	555	carder

Table 6: Top 10 most commonly traded merchandise categories on LC.

Motoyama et al, An Analysis of Underground Forums, 2011

How to make money off financial credentials?

- Money mules
  - Deposits into mules' account from the victim's
  - Mule purchases items using stolen CCN, sells them online
  - Mule withdraws cash from ATMs using victim credentials
- Wires money to (frequently) former Soviet Union

from Richard Hill <hill@hetajobs.com>

subject **Cool Student Job**

to pubs@cs.wisc.edu

reply

Dear Student,

I would like to offer you a new interesting and respectable job!
We are looking for people to work as professional distance-based typists. No experience is needed!
If you're eager to use your skills to make some additional cash, then you might want to consider a home typing position!

All data entry operators work from home and are independent contractors.
You typically set your own hours and work from home on projects that are enjoyable!
Average monthly earnings start from $1000 to $3000 or more.

Requirements:
-Computer with Internet access.
-Good Typing Skills.
-Basic Internet knowledge.
-Basic Computer and Typing Skills.

You will not have to devote full time hours. These assignments can be done on your time.
They may be done in Internet cafes or where ever you can get Internet access!
If you are interested just reply to my email!

Best Regards,

Richard Hill
Local Recruitment Manager

06 Crooks Net Millions in Coordinated ATM Heists

FEB 13

Organized cyber criminals stole almost $11 million in two highly coordinated ATM heists in the final days of 2012, KrebsOnSecurity has learned. The events prompted **Visa** to warn U.S. payment card issuers to be on high-alert for additional ATM cash-out fraud schemes in the New Year.

How to make money off a botnet?

- Rental
  - "Pay me money, and I'll let you use my botnet… no questions asked"
- DDoS extortion
  - "Pay me or I take your legitimate business off web"
- Bulk traffic selling
  - "Pay me to direct bots to websites to boost visit counts"
- Click fraud, SEO
  - "Simulate clicks on advertised links to generate revenue"
  - Cloaking, link farms, etc.
- Theft of monetizable data (eg., financial accounts)
- Data ransom
  - "I've encrypted your harddrive, now pay me money to unencrypt it"
- Advertise products

Botnets

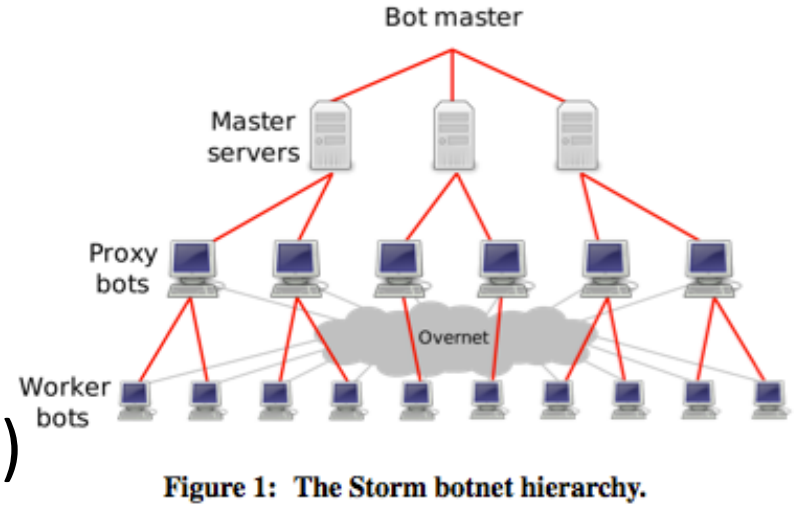

Figure 1: The Storm botnet hierarchy.

- Botnets:
  - Command and Control (C&C)
  - Zombie hosts (bots)
- C&C type:
  - centralized, peer-to-peer
- Infection vector:
  - spam, random/targeted scanning, drive-by exploit
- Usage:
  - What they do: spam, DDoS, SEO, traffic generation, …

Agobot (circa 2002)

- IRC botnet
- Rich feature set:
  - Well-documented, modular codebase
  - IRC-based C&C system
  - Large catalogue of remote exploits
  - Limited code obfuscation and anti-disassembly techniques
  - Built-in data collection
  - Mechanisms to disable antivirus
  - Large set of bot commands

Storm botnet (2007-08)

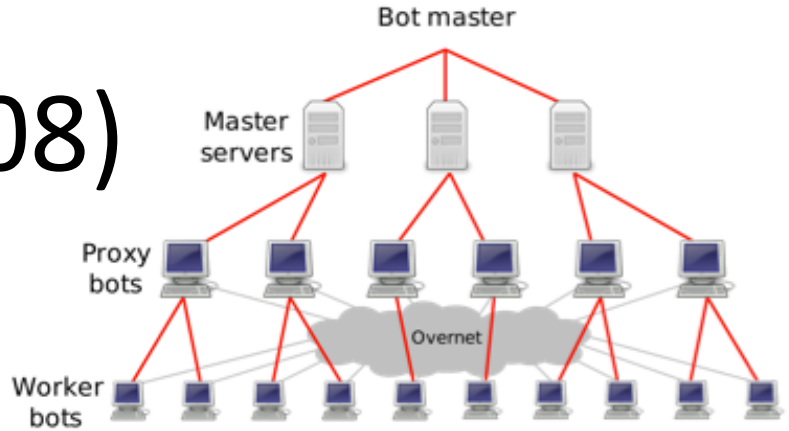

Figure 1: The Storm botnet hierarchy.

- Sept 2007
  - Media: 1 – 50 million bots
  - More likely: 10,000s to 100,000s
- Early spam campaigns used titles such as "230 dead as storm batters Europe."
- Propagated via spam linking to malware
- Thought to be controlled by Russian Business Network

Features:

- Uses P2P (Overnet/Kademlia)
- Uses fast-flux DNS for hosting on named sites
- Binary has gone through many revisions
- Features of P2P network have evolved with time
- Hides on machine with rootkit technology

Enright 2007

See also: http://cseweb.ucsd.edu/~savage/papers/login08.pdf

Overnet DHT

DHT stands for Distributed Hash Table

Peer-to-peer protocol for storing and retrieving data.

Storm uses DHT to store IP addresses of proxies

Proxies maintain contact withBot C&C servers

Publicly addressable bots advertise as able to C&C proxy

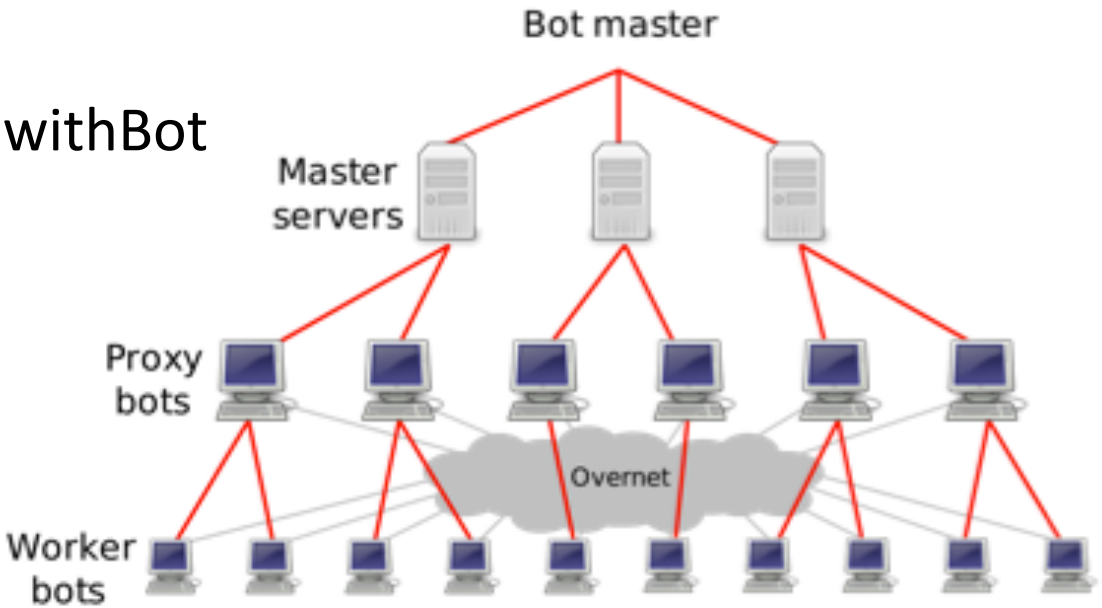

Figure 1: The Storm botnet hierarchy.

Fast-flux DNS

Spam campaign that directs users to ***pharmashop.com***

Single flux:

- Change A record for ***pharmashop.com*** quickly to point to different compromised systems
- Short TTL (e.g., 5 minutes)

Double flux:

- Change NS record for ***pharmashop.com*** to point to different compromised systems

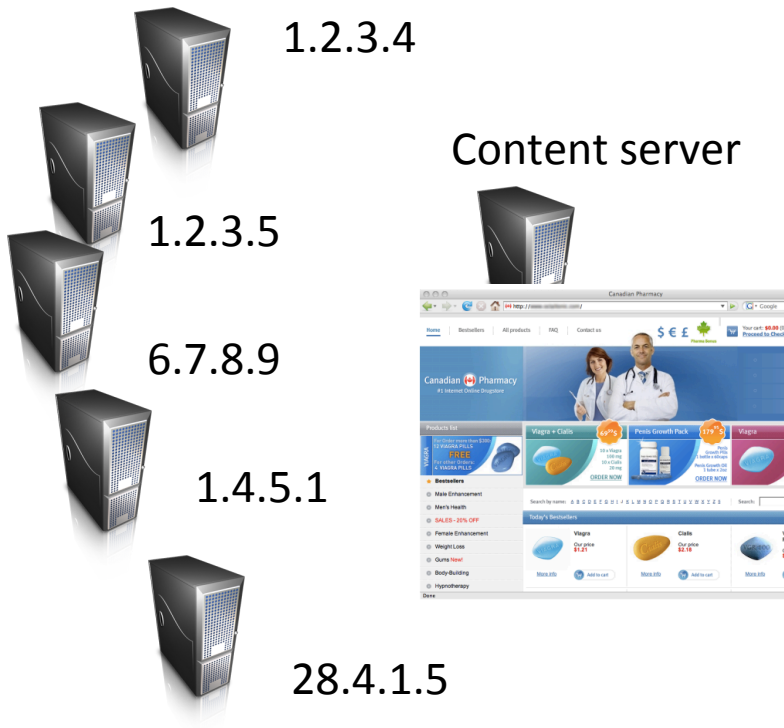

Similar to round-robin DNS as used by major websites

```
;; ANSWER SECTION:
images.pcworld.com.                  900     IN    CNAME   images.pcworld.com.edgesuite.net.
images.pcworld.com.edgesuite.net.  21600     IN    CNAME   a1694.g.akamai.net.
a1694.g.akamai.net.                   20     IN    A       212.201.100.135
a1694.g.akamai.net.                   20     IN    A       212.201.100.141
```

```
;; ANSWER SECTION:
thearmynext.info.       600     IN      A       69.183.26.53
thearmynext.info.       600     IN      A       76.205.234.131
thearmynext.info.       600     IN      A       85.177.96.105
thearmynext.info.       600     IN      A       217.129.178.138
thearmynext.info.       600     IN      A       24.98.252.230
```

```
;; ANSWER SECTION:
thearmynext.info.       600     IN      A       213.47.148.82
thearmynext.info.       600     IN      A       213.91.251.16
thearmynext.info.       600     IN      A       69.183.207.99
thearmynext.info.       600     IN      A       91.148.168.92
thearmynext.info.       600     IN      A       195.38.60.79
```

Holz et al. 2008
http://pi1.informatik.uni-mannheim.de/filepool/publications/fast-flux-ndss08.pdf

Studying fast-flux

CDF of average time between DNS A record address changes

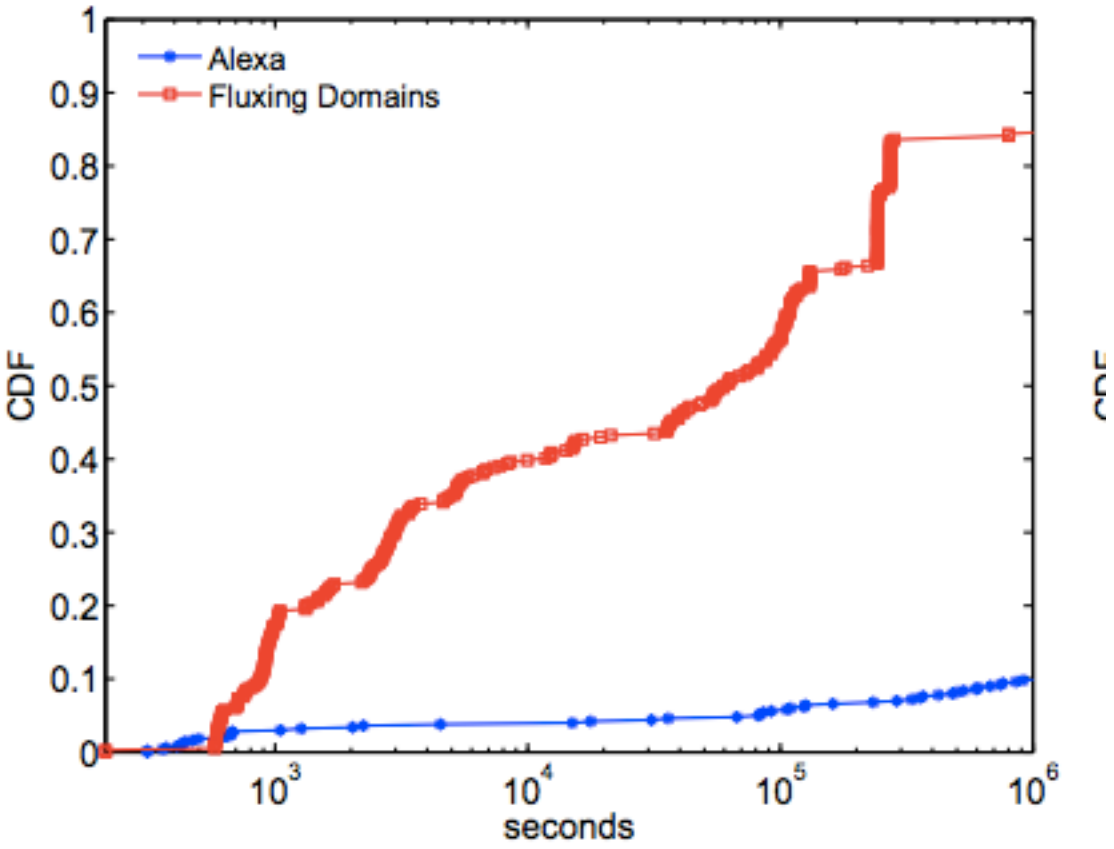

(a) A records

Konte et al. 2008
http://www.cc.gatech.edu/~feamster/papers/fastflux-tr08.pdf

Studying fast-flux

Table 1: Change Rate of Monitored FF Domains

type	Minutes/IP	IP/Day	A-TTL	NS-TTL
average	73.55	55.90	1832.84	37348.75
max	634.50	261.54	21598.03	65535.00
min	5.51	2.27	0	0

Xu et al. 2013
https://media.blackhat.com/us-13/US-13-Xu-New-Trends-in-FastFlux-Networks-WP.pdf

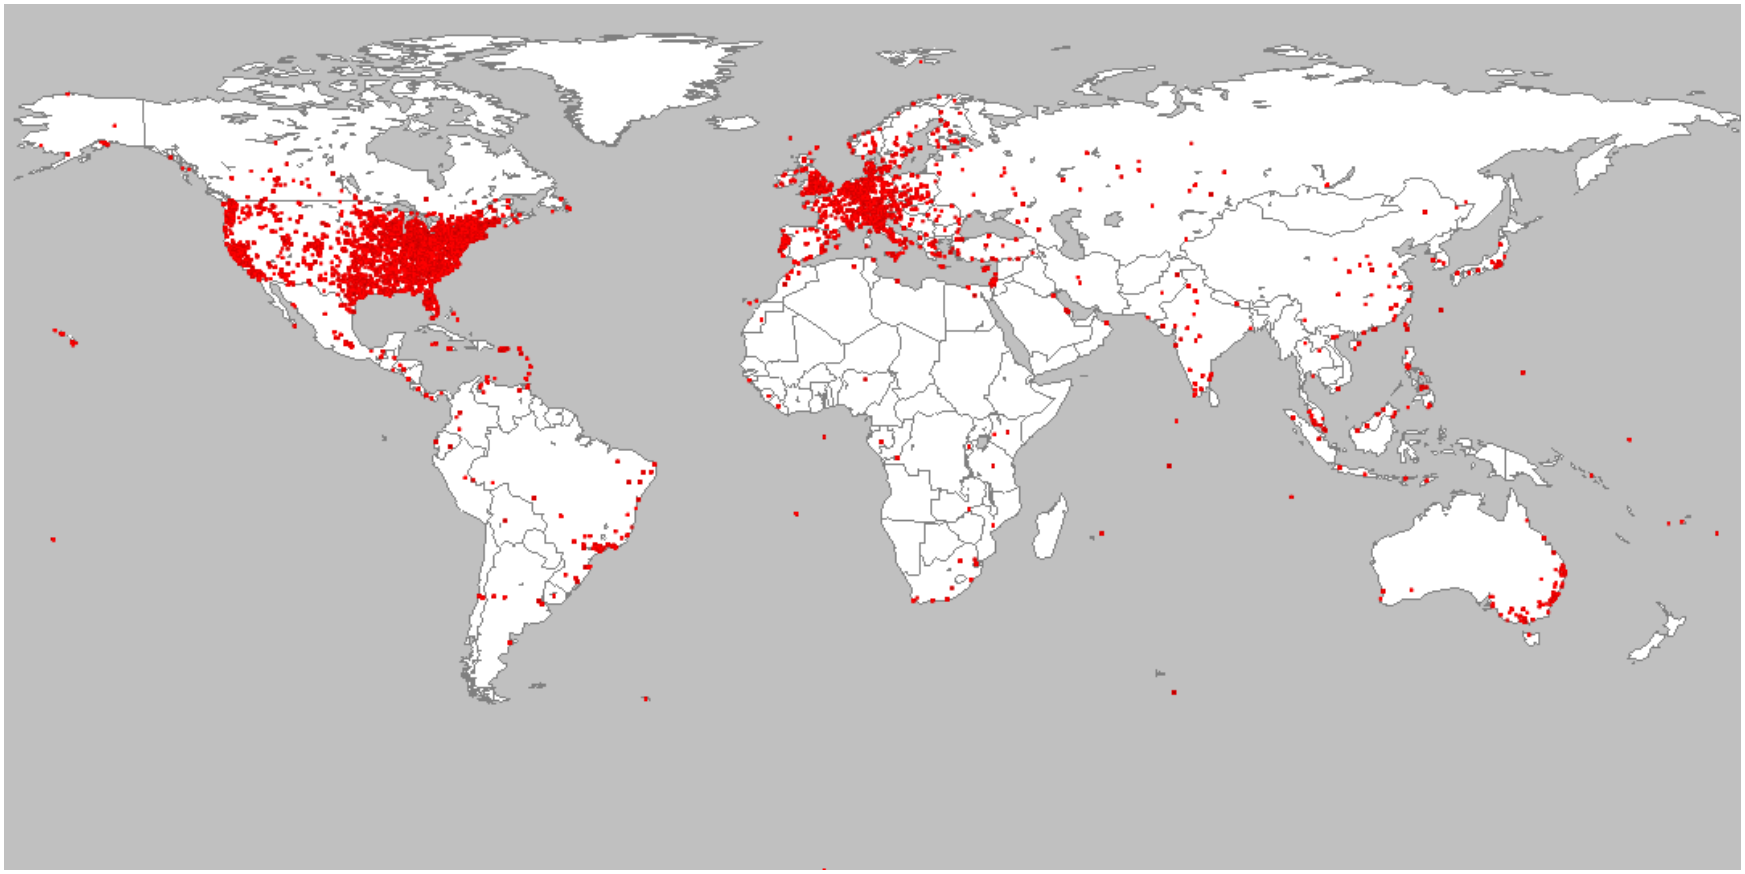

Geolocating bots enumerated for Naguche botnet
Dittrich and Dietrich, “Discovery Techniques for P2P Botnets”

How can we measure botnets?

- Monitor C&C traffic
  - Join IRC channels
  - Network captures at ISPs
- DNS interrogation
  - Identify domain names that point to bots (FastFlux)
- P2P bot crawlers
  - Crawl overlay network to enumerate hosts
- Botnet infiltration
  - Take over C&C
  - Participate as proxies, bots, etc.

Botnet measurement methods

Technique	Description	Pros	Cons
Monitor endpoint	monitor traffic of a bot	simple, generally applicable	limited view, encryption
Internet telescopes	monitor random-scan infection attempts	botnet-wide view	limited applicability
Monitor IRC	record IRC C&C traffic	simple, botnet-wide view	only IRC botnets
DNS redirect	hijack C&C via DNS	measure infection size	limited applicability
Sybil monitoring	monitor numerous bots	simple, passive	resource-intensive, limited view, structured P2P
Botnet crawling	crawl botnet overlay	enumerate large portion of botnet	detectable
DNS cache probing	probe DNS caches for botnet C&C	simple, passive	loose lower-bound
DNSBL counter-intelligence	sniff DNSBL traffic, heuristically identify bots	passive	limited applicability
Flow analysis	detect botnets via flow-based anomaly detection	wide-scale, handles encryption	tailored to IRC botnets

Size estimates from literature as of 2008

Study	Method(s) used	C&C's observed	Largest botnet size		Total # of infected hosts
			infection	effective	
[13]	IRC monitoring	∼100	226,585	–	–
[8]	IRC monitoring	∼180	∼50,000	–	∼300,000
[22]	DNS cache probing	65	–	–	85,000
	IRC monitoring	>100	>15,000	∼3,000	–
[23]	DNS cache probing	100	–	–	88,000
	IRC monitoring	472	∼100,000	>10,000	426,279
[5]	DNS redirection	∼50	>350,000	–	–
[15]	flow analysis	∼376	–	–	∼6,000,000
[7]	botnet crawling	1	∼160,000	∼44,000	–

Figure 2: Size estimates from the literature. All sizes are the maximum ones given in the appropriate study and the final column represents the total number of infected hosts over all botnets encountered.

Size of Storm botnet

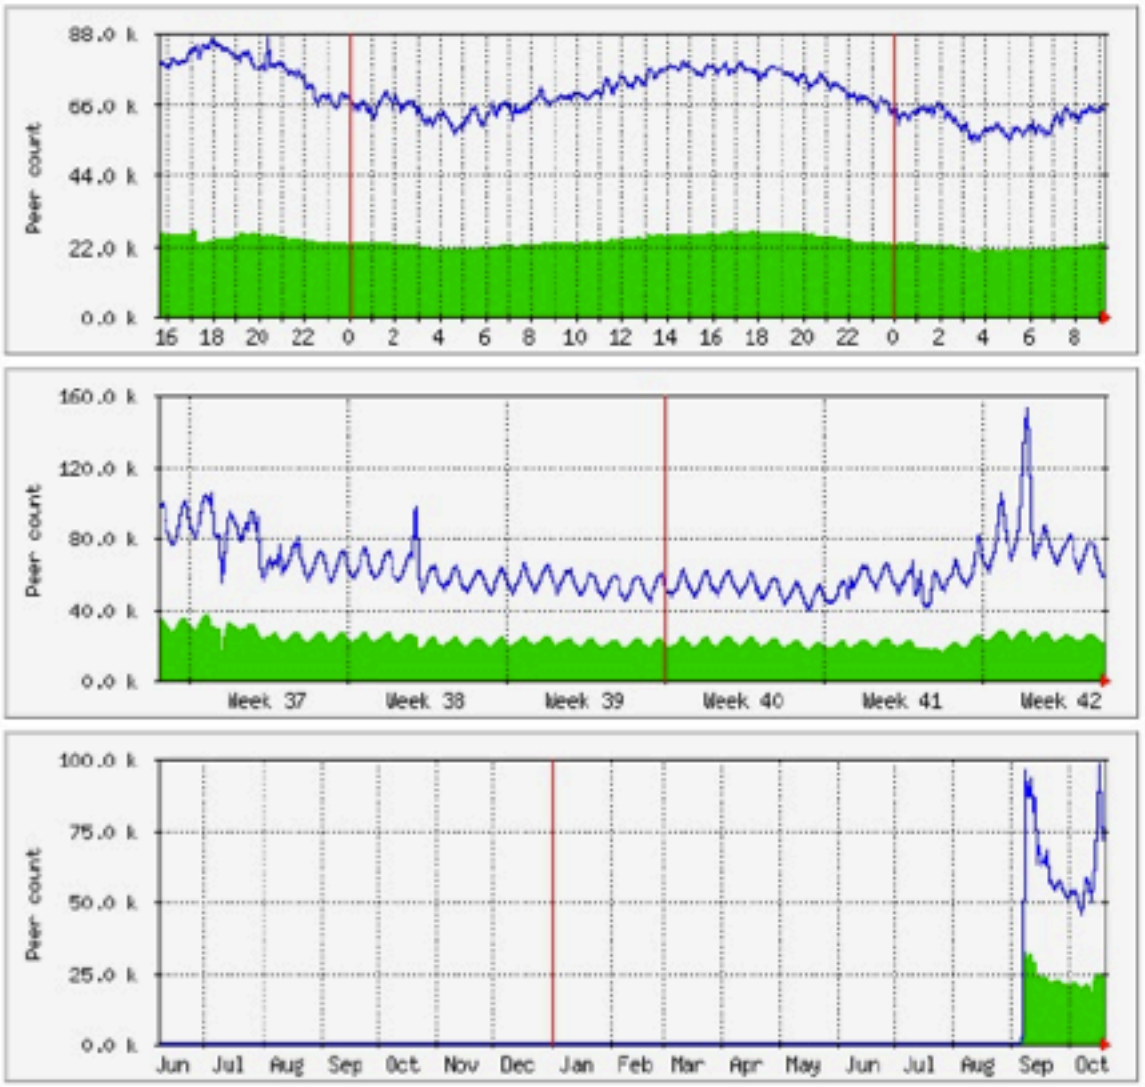

The blue peers count is all peers being probed at a time. This includes live, active, dead, and unknown states. The peers line is not the size of the network. The active line is much closer to the instantaneous size of the network.

It can be seen in the month and year chart that Microsoft made a measurable dent in the network with the MRT Storm (Nuwar) release.

Botnet infiltration studies

- Spamalytics (Kanich et al., 2008)
  - Storm botnet
  - Rewrote spam to redirect to researcher-controlled websites
  - **Goal**: click-through rate measurement

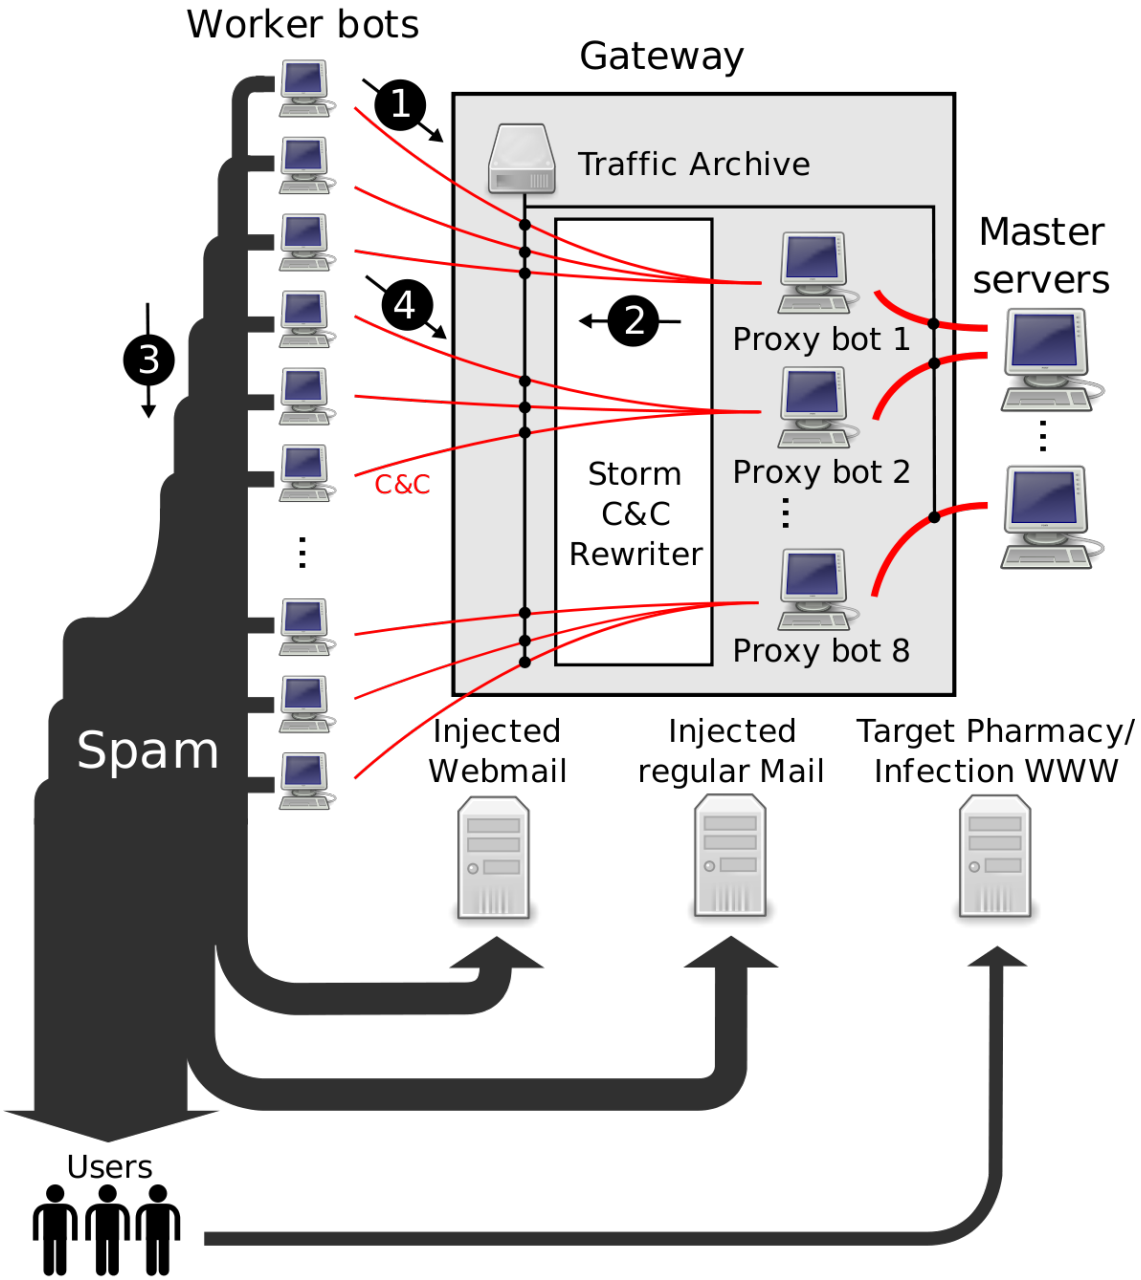

Kanich et al., Spamalytics: An Empirical Analysis of Spam Marketing Conversion, 2008

The victims

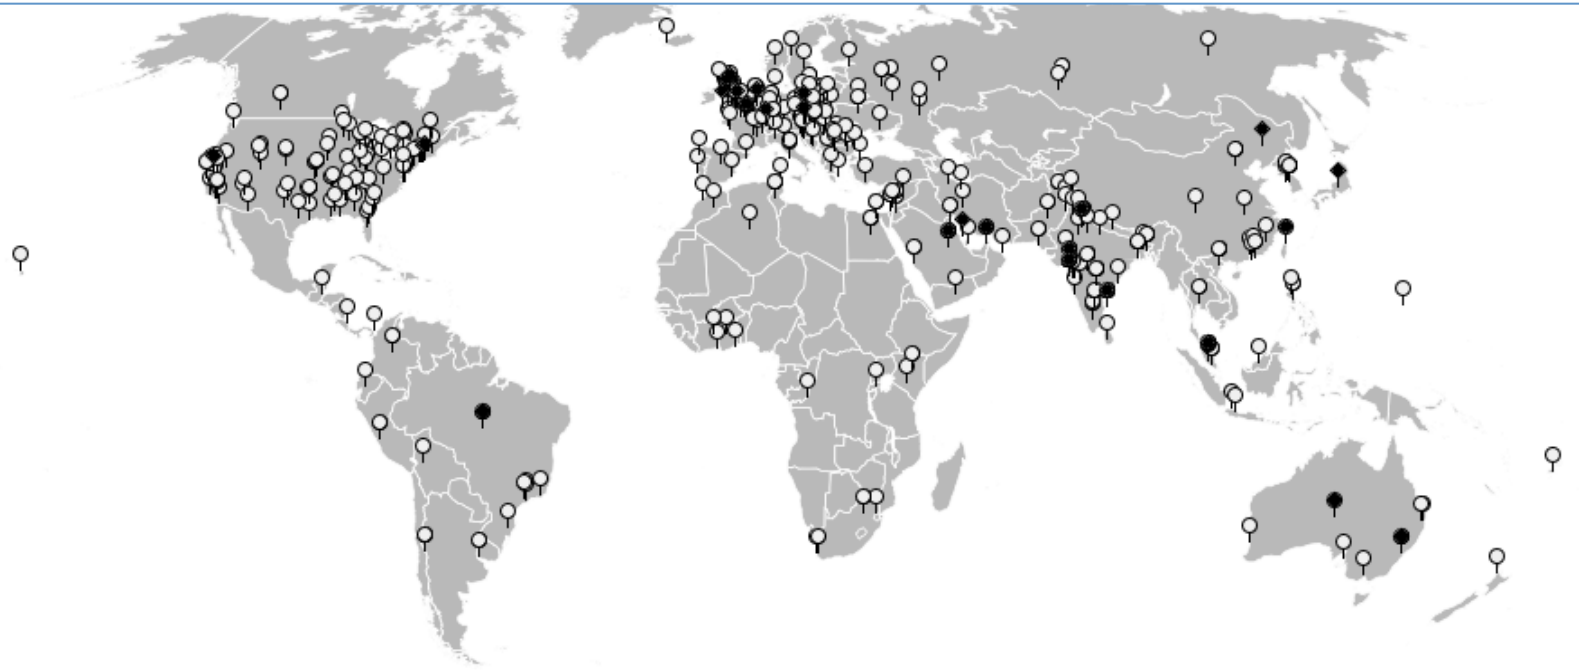

Figure 9: Geographic locations of the hosts that "convert" on spam: the 541 hosts that execute the emulated self-propagation program (light grey), and the 28 hosts that visit the purchase page of the emulated pharmacy site (black).

Kanich et al., Spamalytics: An Empirical Analysis of Spam Marketing Conversion, 2008

Observed Conversion Rate

- 350 million email messages delivered
- 26 day campaign
- 28 “sales”
  - 0.00001%
  - 27 of these male-enhancement products

Botnet takeover studies

- Spamalytics (Kanich et al., 2008)
  - Storm botnet
  - Rewrote spam to redirect to researcher-controlled websites
  - **Goal**: click-through rate measurement
- Torpig C&C sinkholing (Stone-gross et al., 2009)
  - Torpig botnet
  - Setup researcher controlled C&C server (DNS fastflux)
  - **Goal**: analysis of stolen data

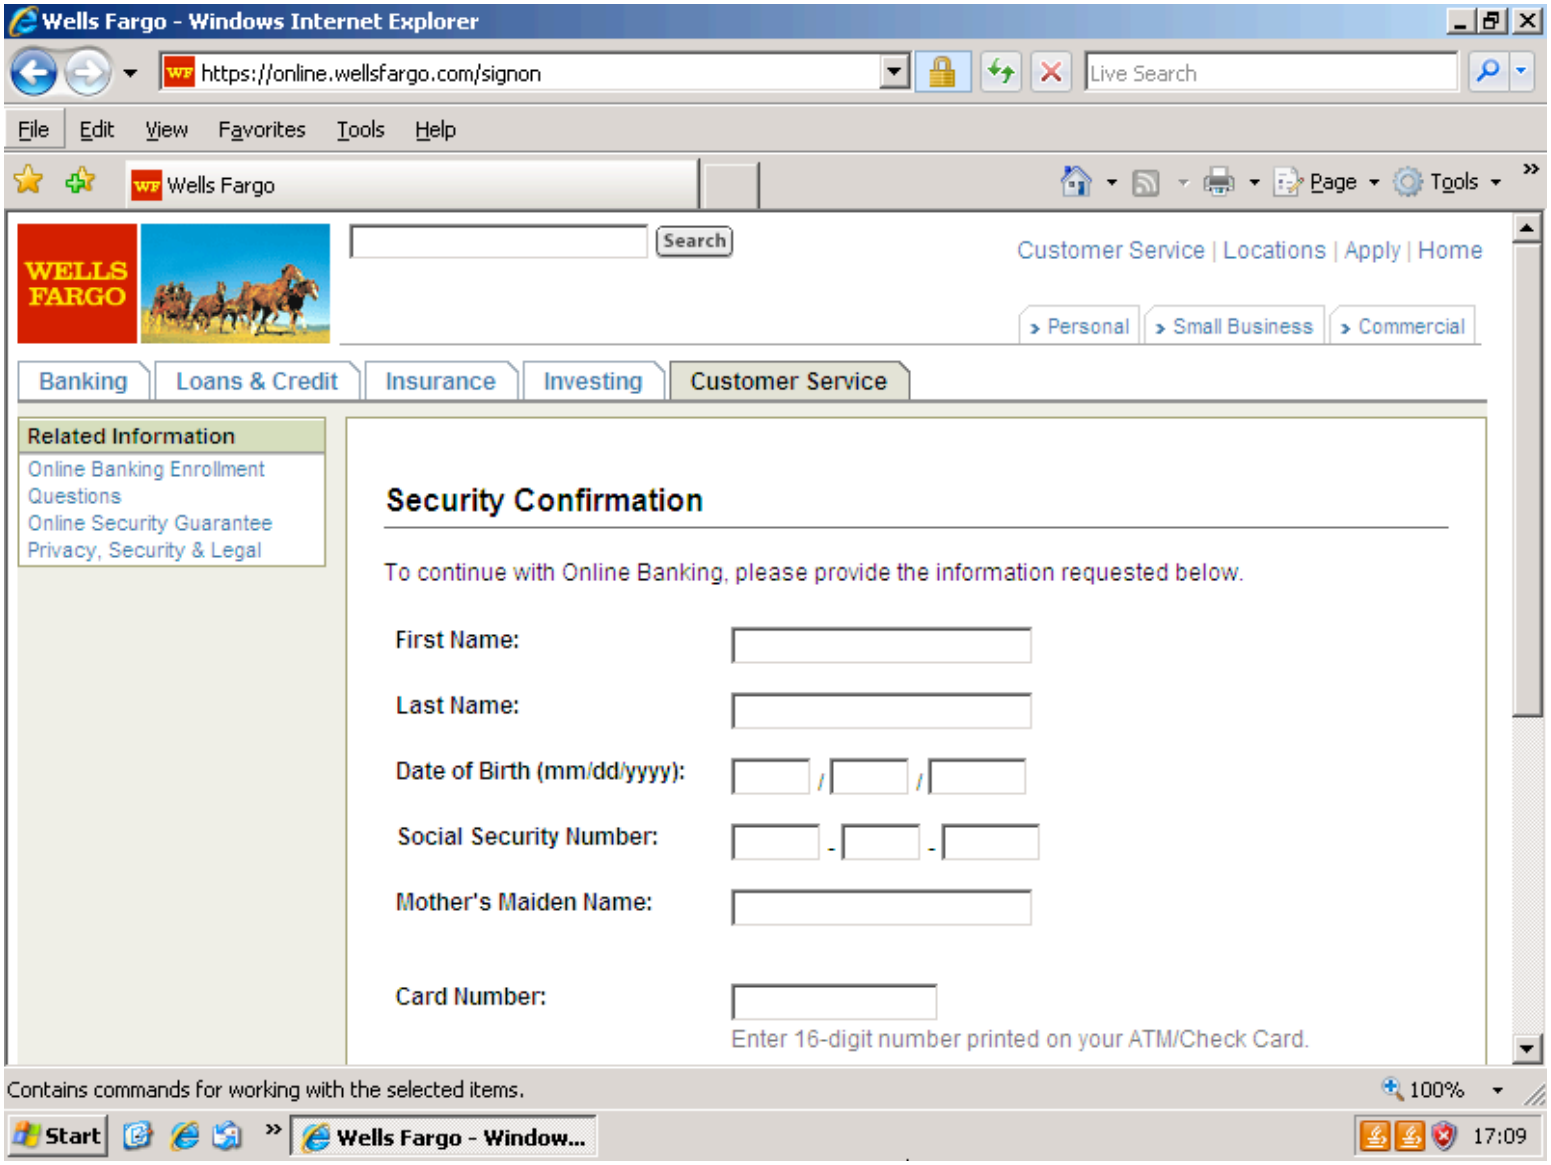

Figure 2: A man-in-the-browser phishing attack.

Stone-Gross et al., Your Botnet is My Botnet: Analysis of a Botnet Takeover, 2009

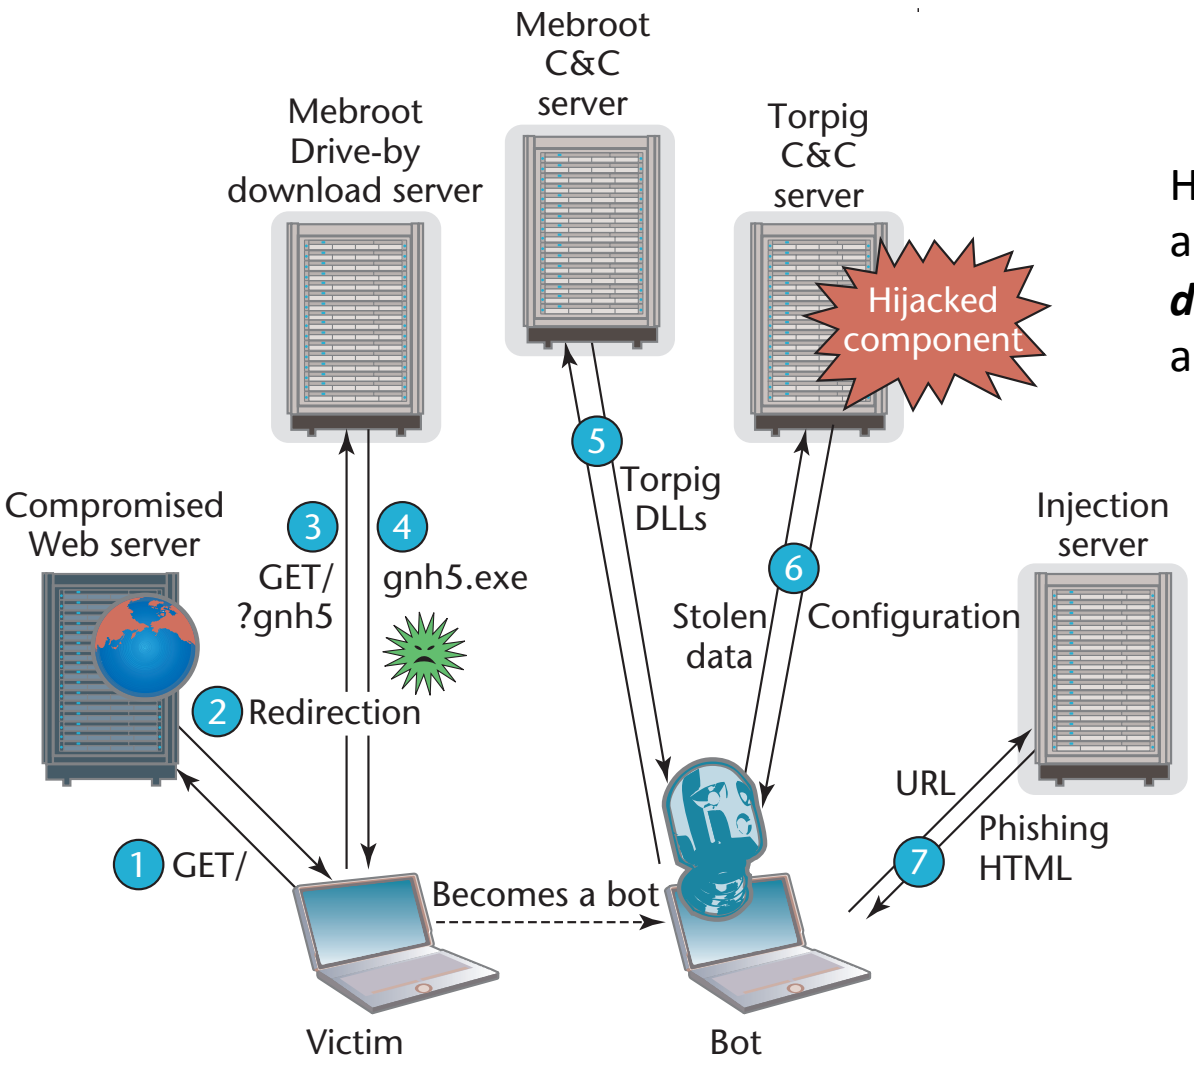

Stone-Gross et al., Your Botnet is My Botnet: Analysis of a Botnet Takeover, 2009

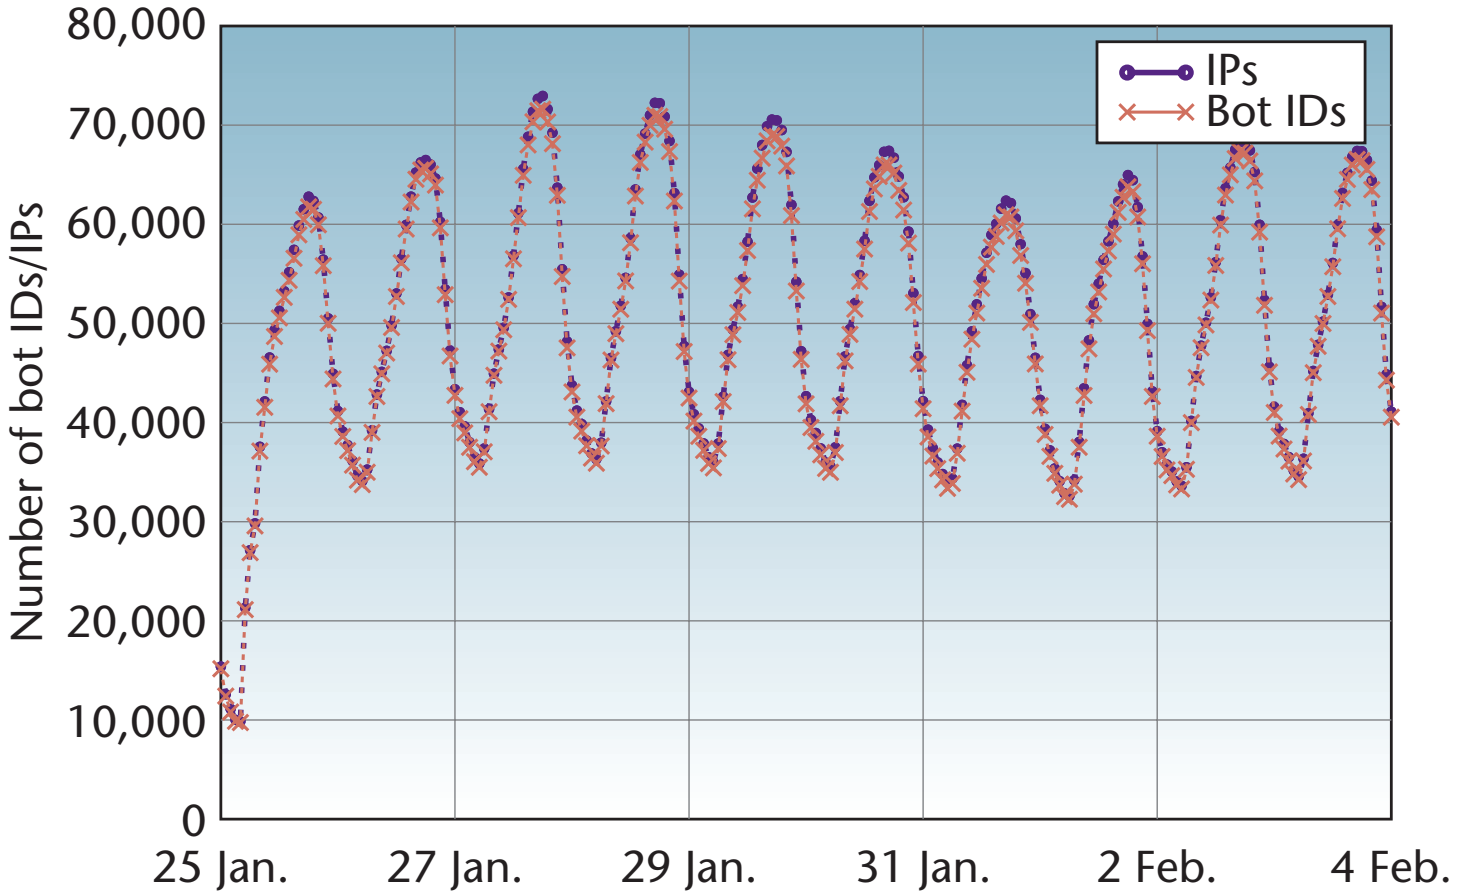

Figure 3. Unique bot IDs and IP addresses per hour. The number of unique IP addresses per hour provides a good estimation of Torpig's live population.

Stone-Gross et al., Your Botnet is My Botnet: Analysis of a Botnet Takeover, 2009

Table 1. Data items sent to our C&C server by Torpig bots.

Data type	Data items
Form data	11,966,532
Email	1,258,862
Windows password	1,235,122
POP account	415,206
HTTP account	411,039
SMTP account	100,472
Mailbox account	54,090
FTP account	12,307

Stone-Gross et al., Your Botnet is My Botnet: Analysis of a Botnet Takeover, 2009

Country	Institutions (#)	Accounts (#)
US	60	4,287
IT	34	1,459
DE	122	641
ES	18	228
PL	14	102
Other	162	1,593
Total	410	8,310

Table 3: Accounts at financial institutions stolen by Torpig.

Stone-Gross et al., Your Botnet is My Botnet: Analysis of a Botnet Takeover, 2009

Botnets

- Botnets:
  - Command and Control (C&C)
  - Zombie hosts (bots)
- C&C type:
  - centralized, peer-to-peer
- Infection vector:
  - spam, random/targeted scanning
- Usage:
  - What they do: spam, DDoS, SEO, traffic generation, …

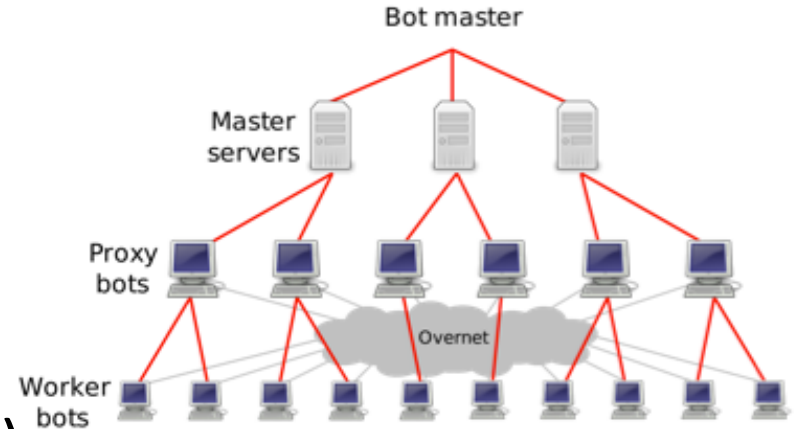

Figure 1: The Storm botnet hierarchy.

Botnet countermeasures?

- Infection prevention
- Infection detection
- C&C take-down
- Undermine the economics
  - Banking take-down

Infection detection & remediation

Anti-Botnet Efforts Still Nascent, But Groups Hopeful

Seven months after a government-industry coalition announced recommendations for ISPs to fight botnets, success is still a long way off

Nov 30, 2012 | 10:06 PM | 0 Comments

By **Robert Lemos, Contributing Writer**
Dark Reading

C&C takedowns

Microsoft Seizes ZeuS Servers in Anti-Botnet Rampage

BY KIM ZETTER 03.26.12 2:45 PM

It's not the first time Microsoft has attempted to take down botnets. The company previously attacked three other botnets — Waledac, Rustock and Kelihos — through similar civil suits that allowed the company to seize web addresses and associated computers. The gains from such takedowns, however, are generally short-lived. After Waledac was targeted, the criminals behind it simply altered their software to thwart easy detection and launched a new botnet.

http://www.wired.com/threatlevel/2012/03/microsoft-botnet-takedown/

Botnet countermeasures?

- Infection prevention
- Infection detection
- C&C take-down
- Undermine the economics
  - Banking take-down

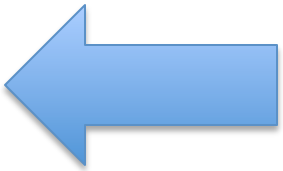

Studying grey/black market products

- Active measurement studies to:
  - Understand (probably illicit) services on web
  - Find ways to defuse underground markets
- Previous studies looked at botnets themselves and victims
- Let's look at the "backend"

Spam-advertised products

- Pharmaceuticals
- Software
- Watches
- etc.

- What is order volume?
- What kinds of things are being purchased?
- What are weak links for disruption?

http://www.rioricopharmacy.com/

From Levchenko et al., “Click Trajectories: End-to-End Analysis of the Spam Value Chain”, IEEE Symposium on Security and Privacy, 2011

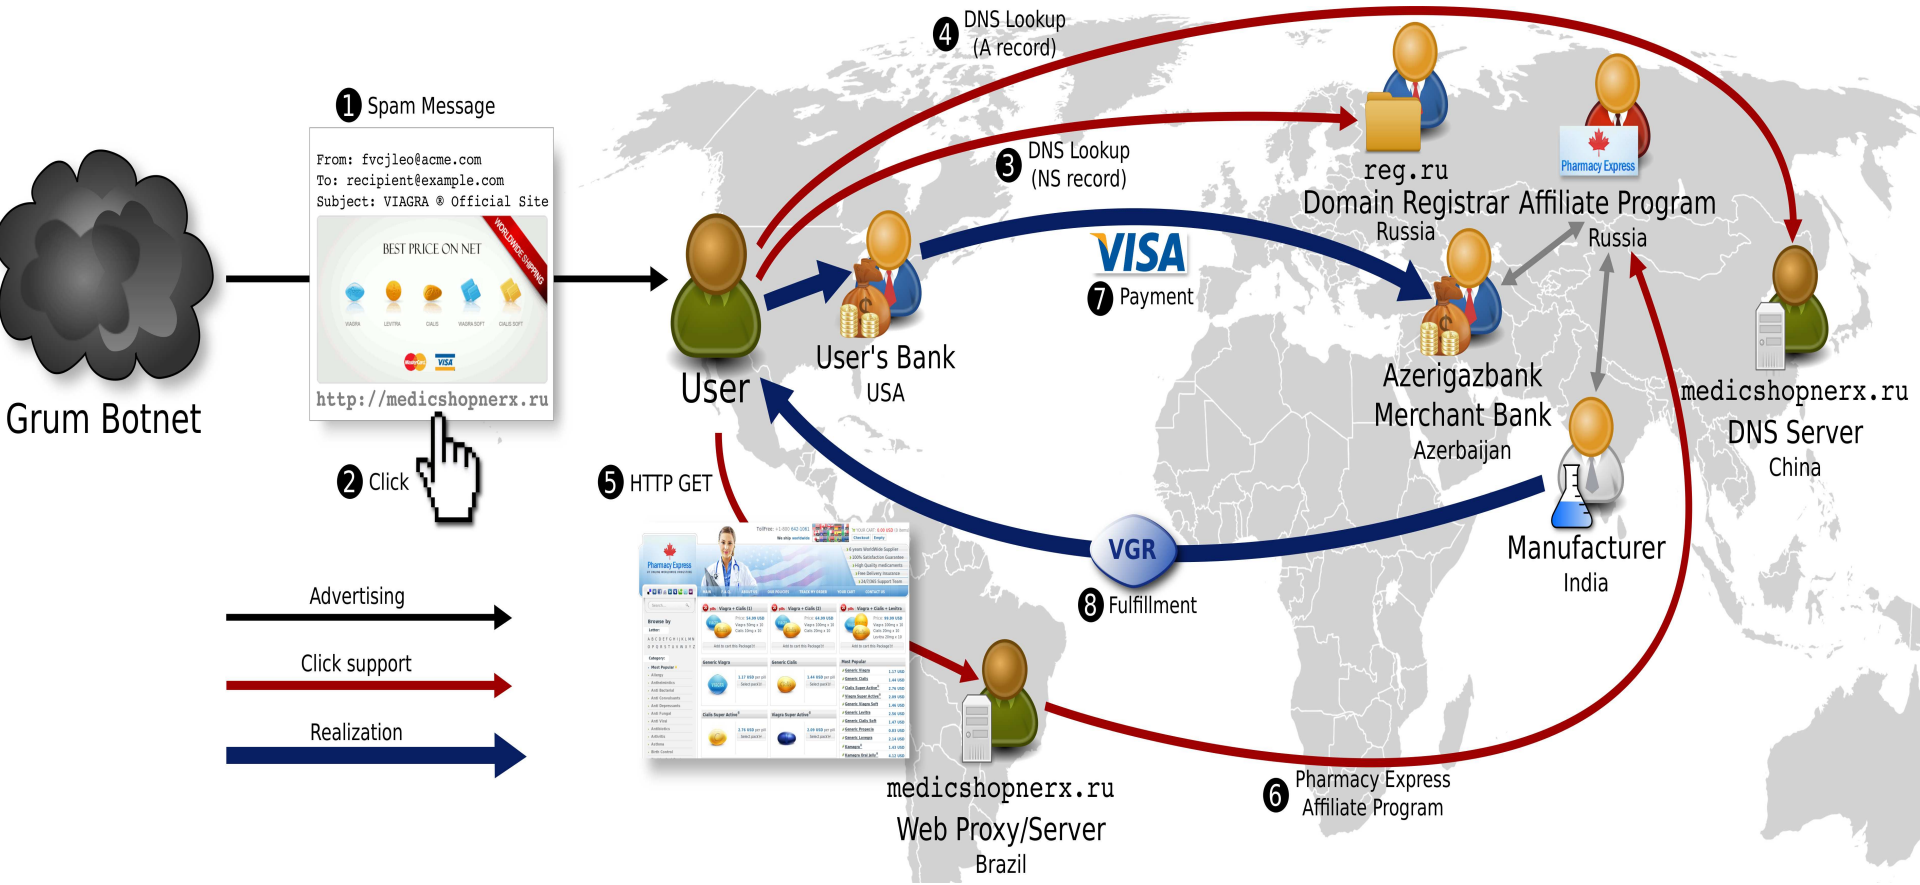

Measurement apparatus #1

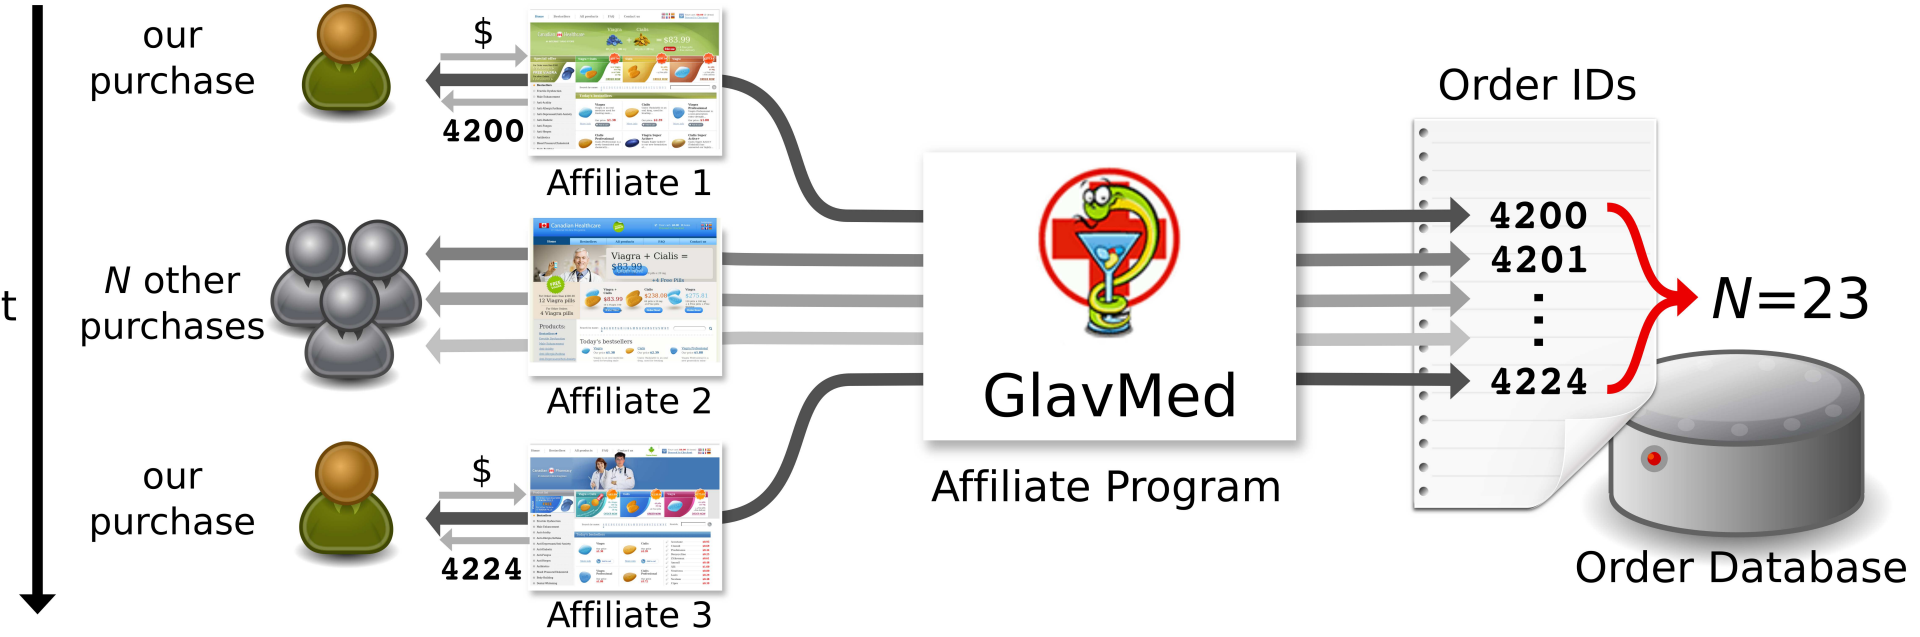

Kanich et al., Show Me the Money: Characterizing Spam-advertised Revenue, 2011

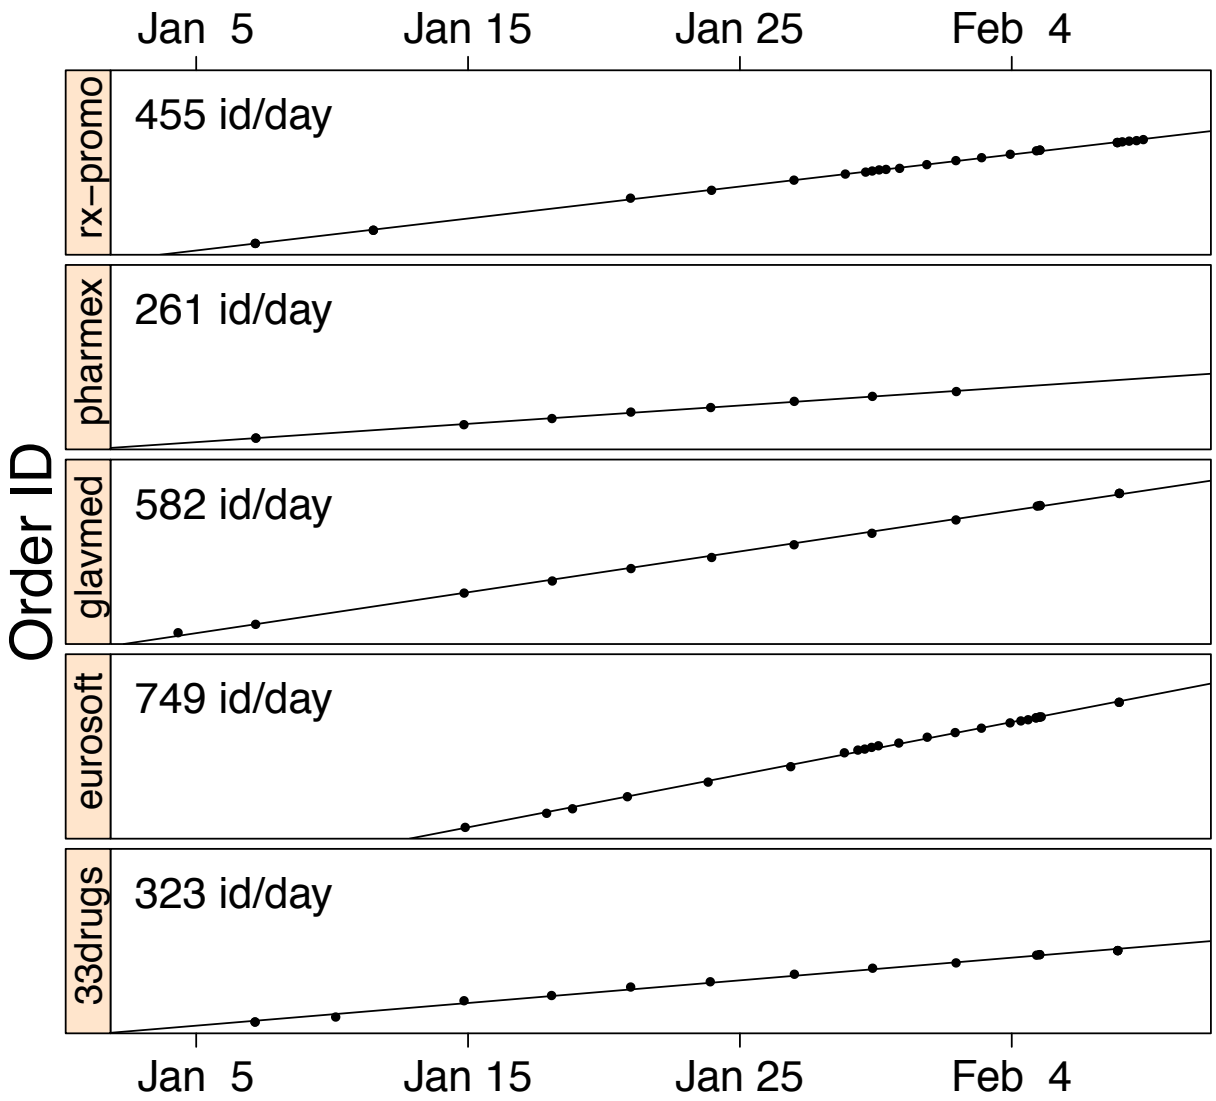

Kanich et al., Show Me the Money: Characterizing Spam-advertised Revenue, 2011

Measurement Apparatus #2

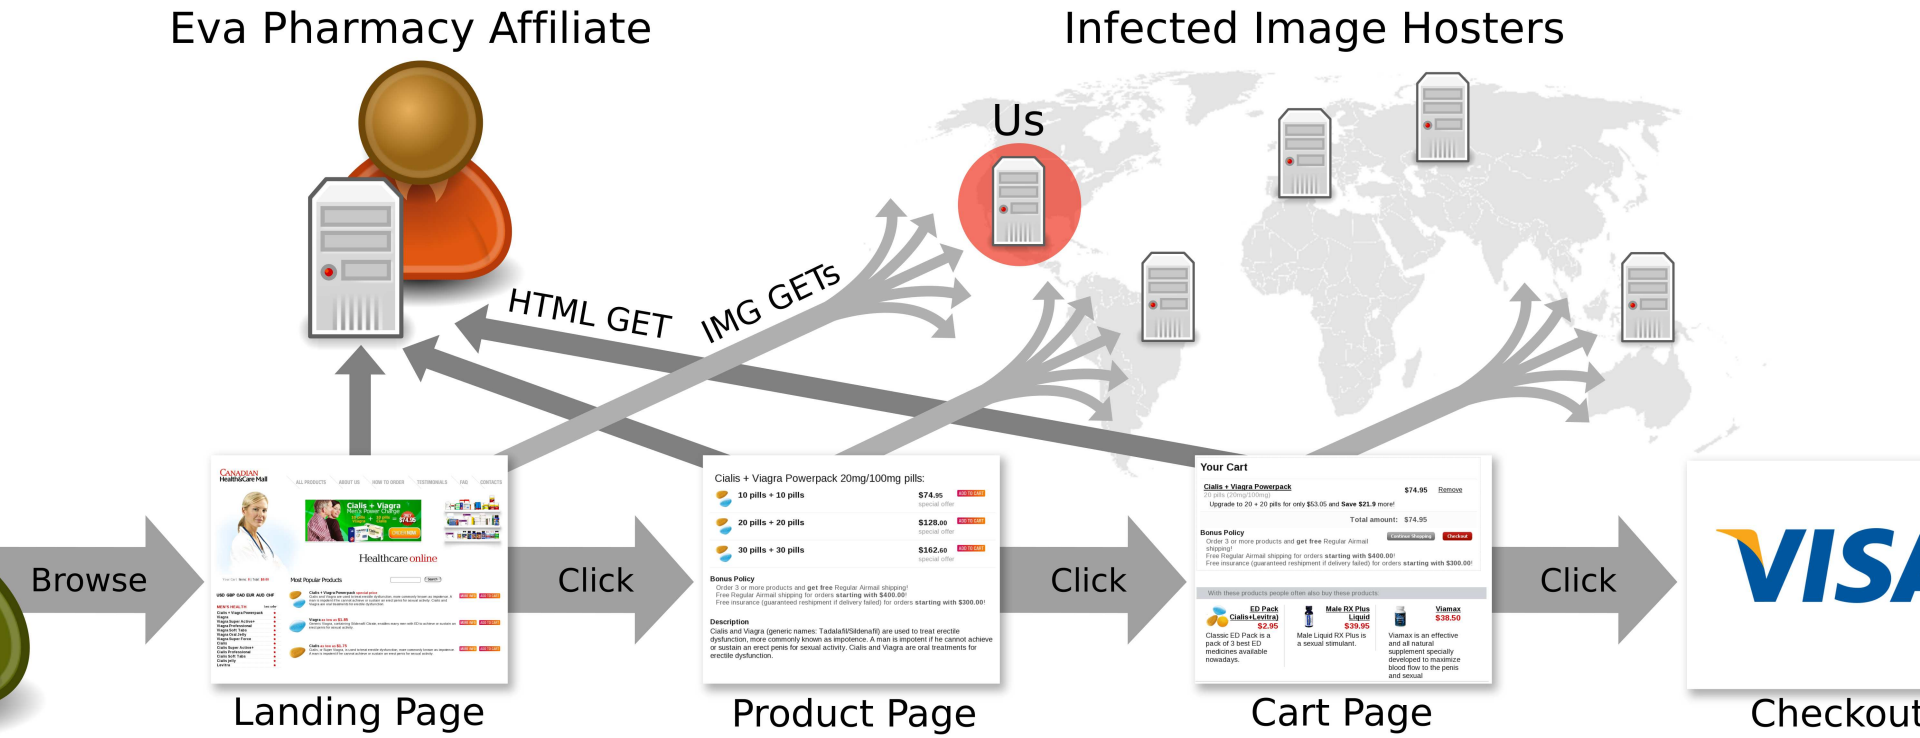

Product	Quantity	Min order
Generic Viagra	568	$78.80
Cialis	286	$78.00
Cialis/Viagra Combo Pack	172	$74.95
Viagra Super Active+	121	$134.80
Female (pink) Viagra	119	$44.00
Human Growth Hormone	104	$83.95
Soma (Carisoprodol)	99	$94.80
Viagra Professional	87	$139.80
Levitra	83	$100.80
Viagra Super Force	81	$88.80
Cialis Super Active+	72	$172.80
Amoxicillin	47	$35.40
Lipitor	38	$14.40
Ultram	38	$45.60
Tramadol	36	$82.80
Prozac	35	$19.50
Cialis Professional	33	$176.00
Retin A	31	$47.85

Kanich et al., Show Me the Money: Characterizing Spam-advertised Revenue, 2011

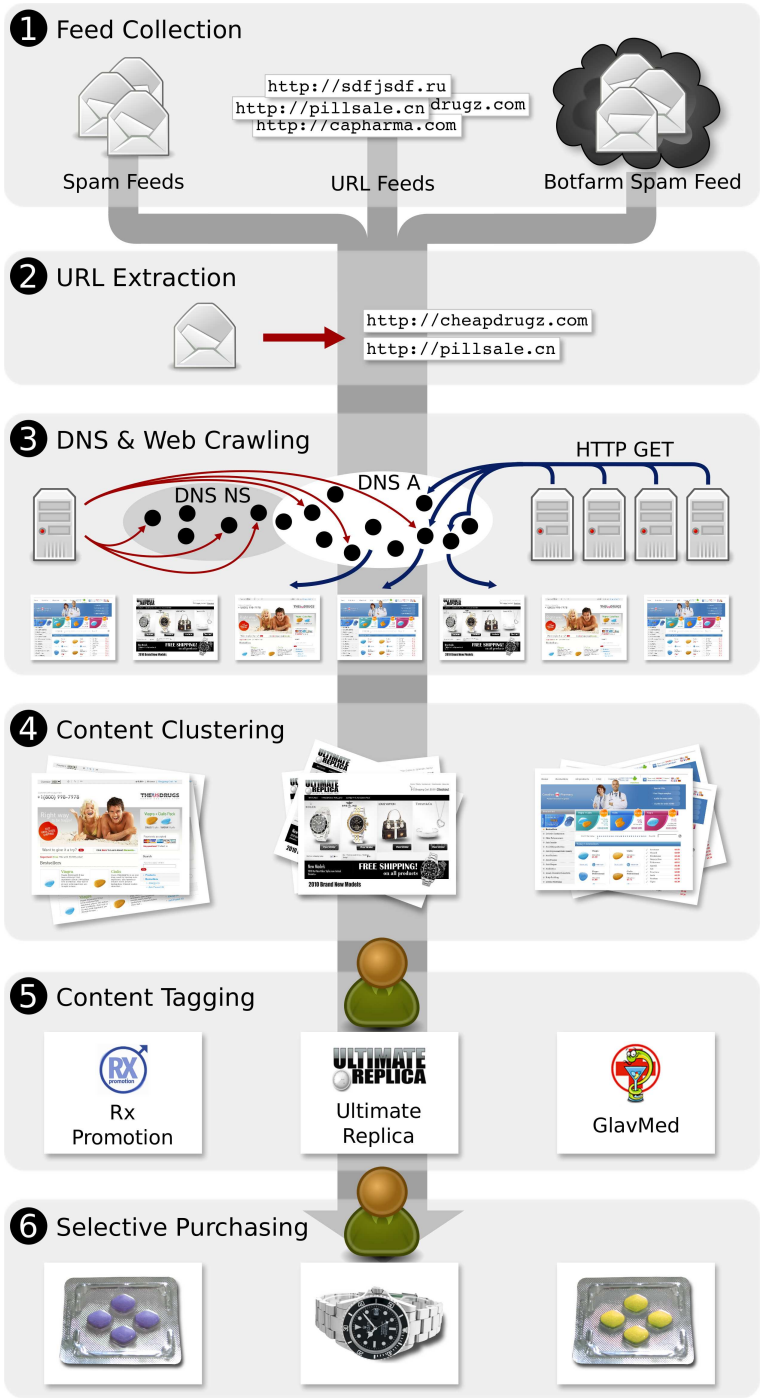

Levchenko et al., Click Trajectories: An End-to-End Analysis of the Spam Value Chain, 2011

- 120 items purchased
- 76 authorized
- 56 settled
- 49 products delivered

- 2 sent after mailbox lease ended
- 2 no follow-up email
- 2 resent after mailbox lease ended
- 1 promised refund (never obtained)

Supplier	*Item*	*Origin*	*Affiliate Programs*
Aracoma Drug	Orange bottle of tablets (pharma)	WV, USA	ClFr
Combitic Global Caplet Pvt. Ltd.	Blister-packed tablets (pharma)	Delhi, India	GlvMd
M.K. Choudhary	Blister-packed tablets (pharma)	Thane, India	OLPh
PPW	Blister-packed tablets (pharma)	Chennai, India	PhEx, Stmul, Trust, ClFr
K. Sekar	Blister-packed tablets (pharma)	Villupuram, India	WldPh
Rhine Inc.	Blister-packed tablets (pharma)	Thane, India	RxPrm, DrgRev
Supreme Suppliers	Blister-packed tablets (pharma)	Mumbai, India	Eva
Chen Hua	Small white plastic bottles (herbal)	Jiangmen, China	Stud
Etech Media Ltd	Novelty-sized supplement (herbal)	Christchurch, NZ	Staln
Herbal Health Fulfillment Warehouse	White plastic bottle (herbal)	MA, USA	Eva
MK Sales	White plastic bottle (herbal)	WA, USA	GlvMd
Riverton, Utah shipper	White plastic bottle (herbal)	UT, USA	DrMax, Grow
Guo Zhonglei	Foam-wrapped replica watch	Baoding, China	Dstn, UltRp

Table VI: List of product suppliers and associated affiliate programs and/or store brands.

Levchenko et al., Click Trajectories: An End-to-End Analysis of the Spam Value Chain, 2011

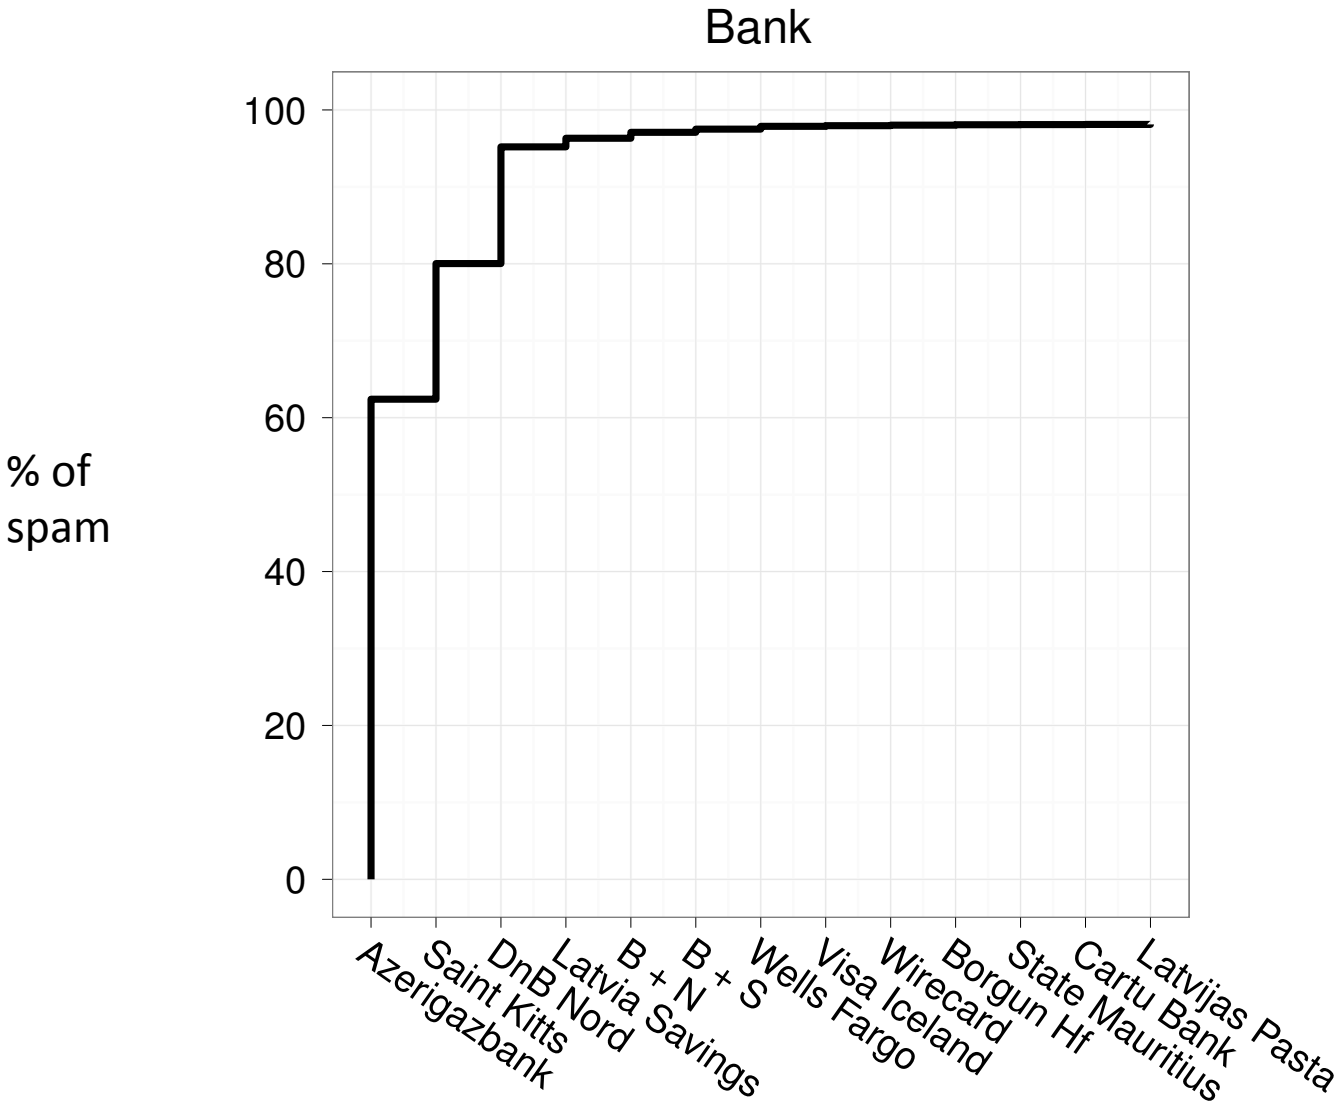

Levchenko et al., Click Trajectories: An End-to-End Analysis of the Spam Value Chain, 2011

Can we throttle abuse by targeting merchant accounts at banks?

- McCoy et al., Priceless: The Role of Payments in Abuse-advertised Goods, 2012
- Made purchases to pharma and software OEM programs, while also working with brandholders to make complaints to Visa/MC

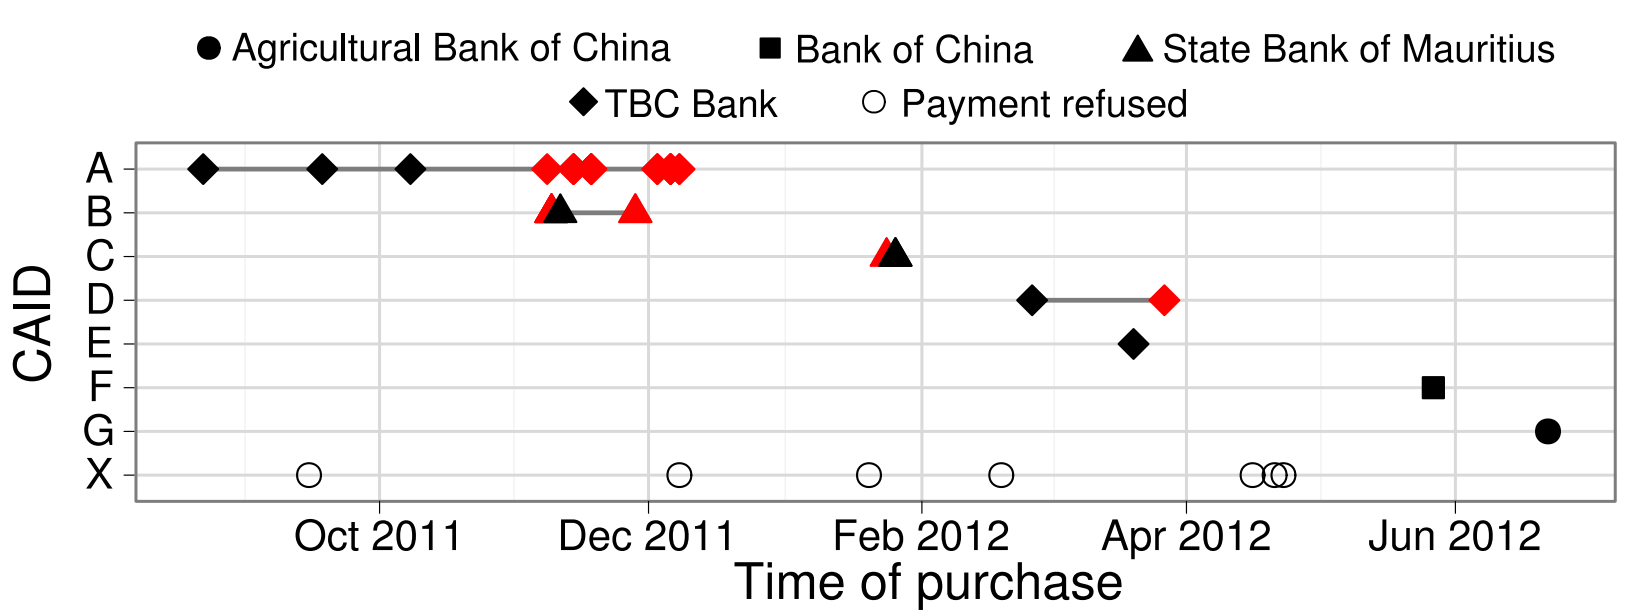

Figure 4: Example of a program receiving complaints to a card network. Rows denote distinct merchant descriptors; row "X" shows refused orders.

Wrote one eloquent affiliate in March of this year, "Right now most affiliate eprograms have a mass of declines, cancels and pendings, and it doesn't depend much on the program IMHO, there is a general sad picture, f!$#@ing Visa is burning us with napalm."

McCoy et al., Priceless: The Role of Payments in Abuse-advertised Goods, 2012

Ethics

- We ha
  - mea
  - part
  - takir
  - purc
  - port scanning victims
- Ethics discussion in papers:
  - short discussion justifying lack of harm
  - “beyond the scope of this work”

From paper on Torpig takeover (Stone-Gross et al.)

PRINCIPLE 1. The sinkholed botnet should be operated so that any harm and/or damage to victims and targets of attacks would be minimized.

PRINCIPLE 2. The sinkholed botnet should collect enough information to enable notification and remediation of affected parties.

Traffic sellers

- Click fraud
- Click traffic sellers
  - grey-market
  - Class project pilot study to see what these sellers are all about
    - Botnet traffic?
    - Legitimate project?
  - http://cseweb.ucsd.edu/~tristenp/buytraffic/

REVISITORS.COM
Targeted Traffic Since 2005

TOLL FREE:
877.389.3330
101 California St - San Francisco, CA

McAfee SECURE
TESTED 08-DEC

You can't make sales if don't have VISITORS

"30 days unlimited traffic"

Stop getting scammed from traffic sellers!
This is real quality traffic that
We use for own sites.

INCREASE WEB TRAFFIC
GUARANTEED!

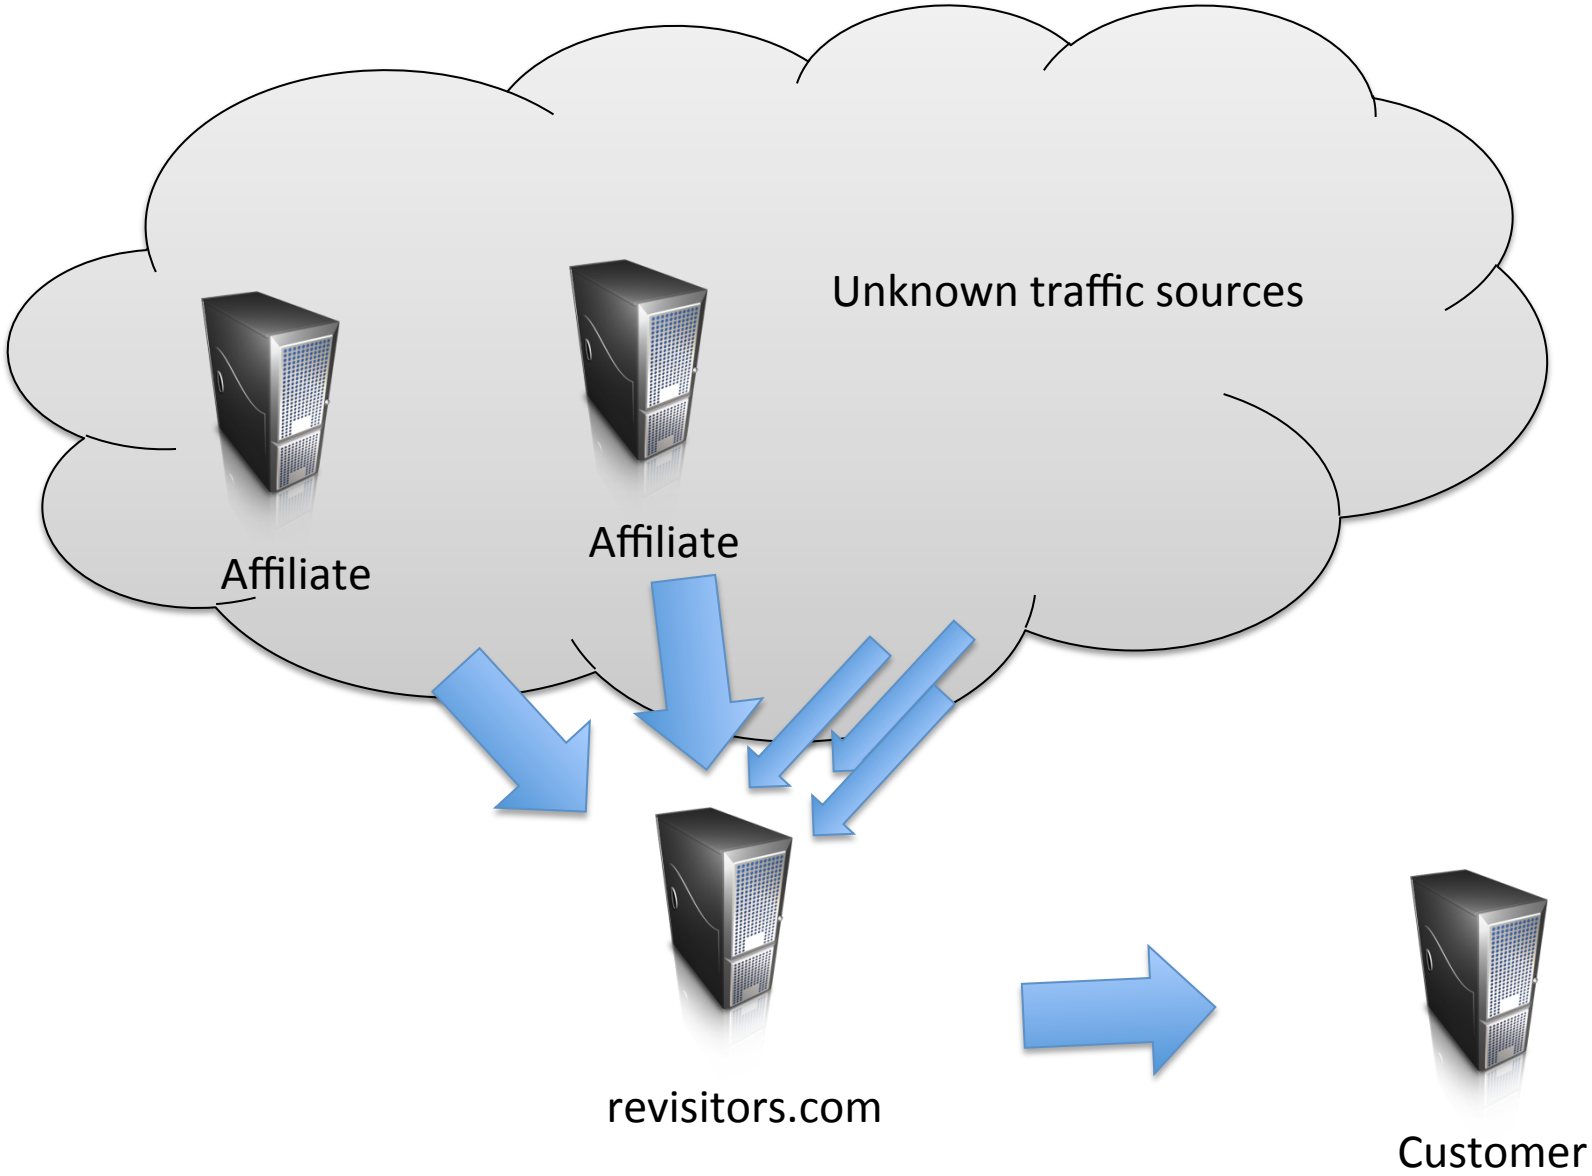
Unknown traffic sources
Affiliate
Affiliate
Affiliate
revisitors.com
Customer

Click traffic sellers

Quality of website's English ↓

Web site	CP10k	Claimed traffic source
www.trafficdeliver.com	~$34.69	"Advertiser exchange"
revisitors.com	~$48.95	Recently expired domain redirection?
qualitytrafficsupply.com	~$55.00	Contextual advertisements
mediatraffic.com	~$70	AdWare (Voomba) pop-ups

Targeted vs. untargeted: specify geographic preferences

Affiliate networks: paid to send traffic

Traffic resellers: resell purchased traffic

Experimental methodology

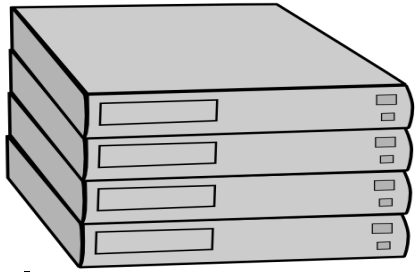

(1) Setup several web sites (xxx.sysnet.ucsd.edu)

2 pages: index.html is landing site
lucky.html linked to by index.html

Example site linked from
webpage

(2) Attempt to purchase web traffic

Used temporary VISA number, but real name, etc.

(3) Sit back and let the research data come to us ...

Adventures in purchasing web traffic...

Giving people money not as easy as I expected:

RE: Refund - [2423-DLXC-4301] [82a2e44b]

2Checkout Help Desk	========= Please enter your reply ABOVE above this line ========= Hello Tom, ...	Dec 6 (5 days ago)
2Checkout Help Desk	A staff member has replied to your question: Seasons Greetings Tom, Thank you...	Dec 6 (5 days ago)
2Checkout Help Desk	Thank you for adding a message to your question. We will respond to your mess...	Dec 6 (4 days ago)

2Checkout Help Desk to me — show details Dec 6 (4 days ago) — Reply

========= Please enter your reply ABOVE above this line =========

Hello Tom,

A staff member has replied to your question:

Dear Tom,

Thank you for contacting 2Checkout.com. I apologize for the delay in responding to your inquiry. The order was actually canceled trafficdeliver.com. They believe the order to be fraudulent. I have forwarded your inquiry to trafficdeliver.com. They will be contacting you via e-mail shortly. If you do not receive a response in a timely manner, please feel free to reopen this ticket for additional assistance.

Looking to make your holidays happier? 2Checkout makes it easy! Simply visit your favorite search engine and type in 2Checkout + and the type of merchandise you are looking for. It's the easy way to enjoy a fast, safe shopping experience online.

Thank You,
Josh Karamian
Customer Care
2Checkout.com
http://www.2Checkout.com

When did traffic arrive?

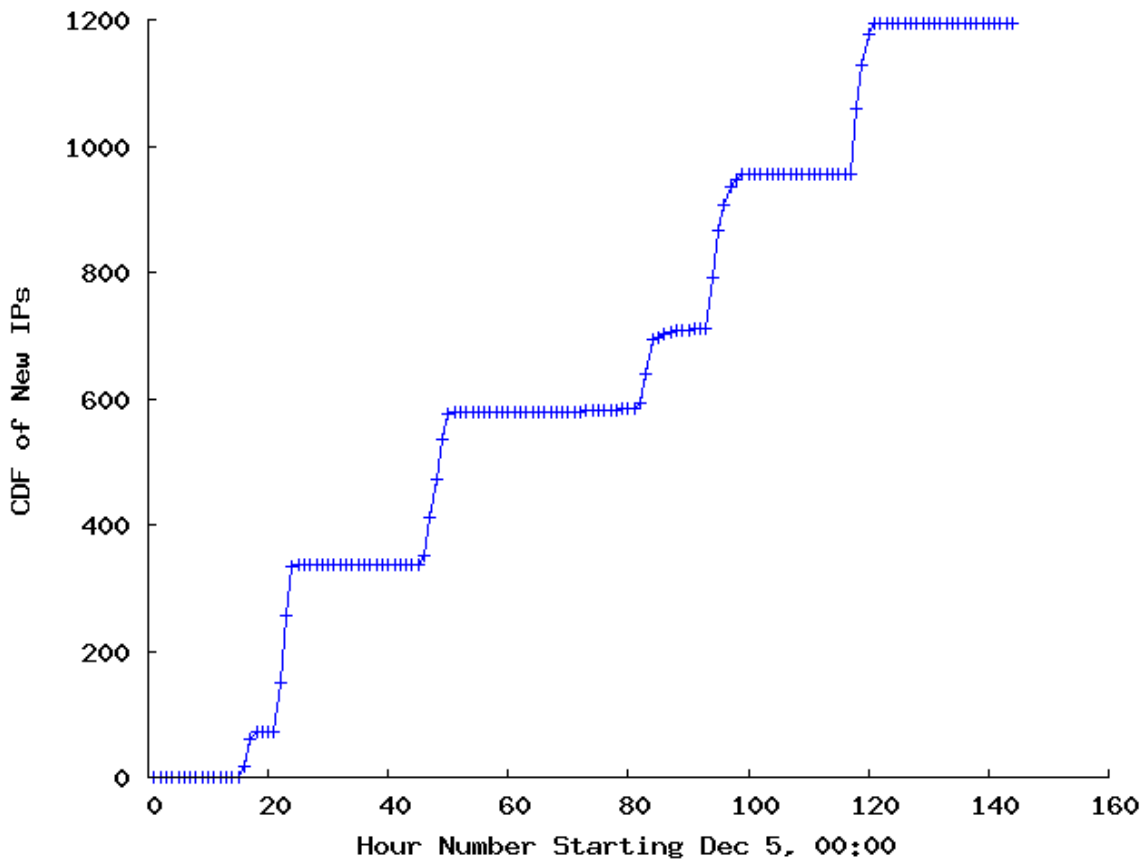

When did traffic arrive?

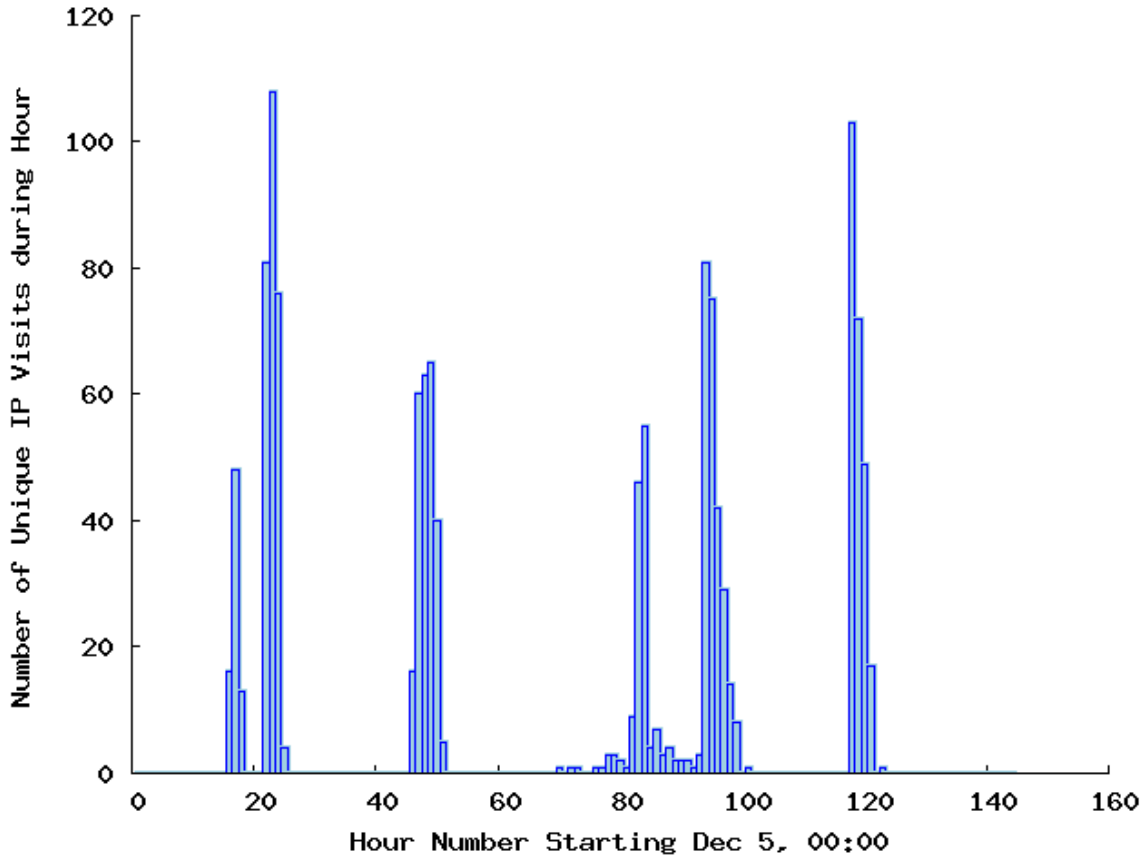

- Not a typical pattern for traffic

When did traffic arrive?

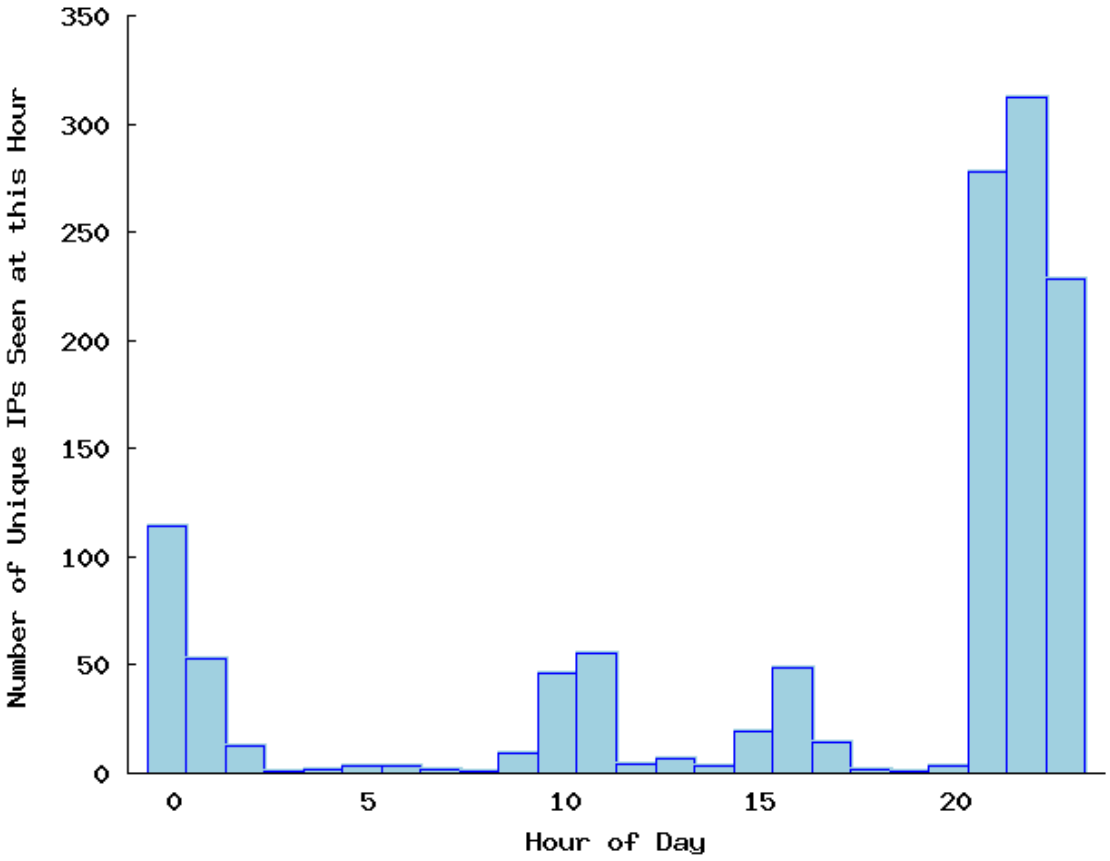

- Traffic has really high-degree of temporal proximity
- Anecdote: many IPs visit times clustered within seconds

Is the traffic from bots or other malware?

Source	Num IPs	Percentage
CBL	21	1.7%
Current Storm	0	0.0%

Other interesting anecdotal evidence

4 HEAD requests from distinct IPs with referrer

http://www.routetraffic.net/delivery/statistics/
8x0ada67md29fk799sa4.html

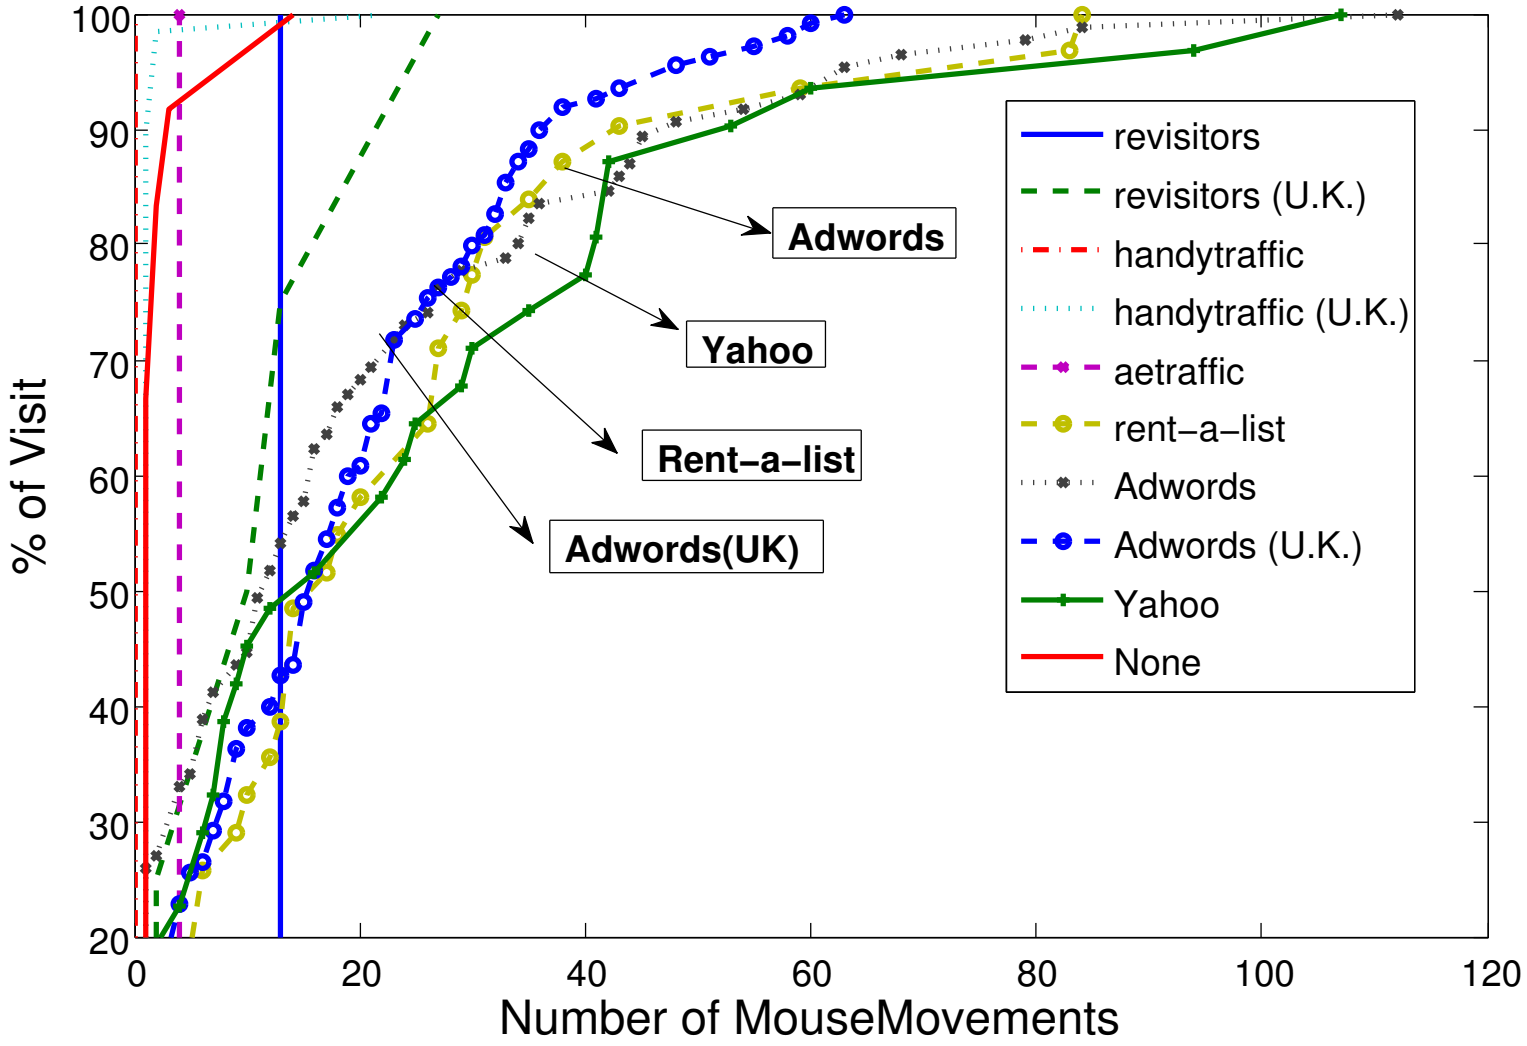

(b) CDF of # of mouse moves per visit across all visits

Zhang et al., Got Traffic? An Evaluation of Click Traffic Providers, 2011

-------- Original Message --------
Subject: Email Alert From UW-Madison Computer Sciences
Date: Wed, 5 Dec 2012 12:49:27 -0430 (VET)
From: cs.wisc.edu <asantanap@cantv.net>
To: <tannenba@cs.wisc.edu>, <swright@cs.wisc.edu>, <swift@cs.wisc.edu>, <sweep@cs.wisc.edu>, <sumit@cs.wisc.edu>, <suman@cs.wisc.edu>, <suhui@cs.wisc.edu>, <subbarao@cs.wisc.edu>, <suan@cs.wisc.edu>, <stuart@cs.wisc.edu>, <strik@cs.wisc.edu>, <street@cs.wisc.edu>, <stever@cs.wisc.edu>, <stefanic@cs.wisc.edu>, <srour@cs.wisc.edu>, <sriram@cs.wisc.edu>, <srikris@cs.wisc.edu>, <sray@cs.wisc.edu>, <soni@cs.wisc.edu>, <solomon@cs.wisc.edu>, <sohi@cs.wisc.edu>, <soc-culture-greek-request@cs.wisc.edu>, <smurphy@cs.wisc.edu>, <smoler@cs.wisc.edu>, <skrentny@cs.wisc.edu>, <sklein@cs.wisc.edu>, <sjha@cs.wisc.edu>, <sigarch-members@cs.wisc.edu>, <shukla@cs.wisc.edu>, <shuchi@cs.wisc.edu>, <shoup@cs.wisc.edu>, <shiliang@cs.wisc.edu>, <shavlikq@cs.wisc.edu>, <shavlik@cs.wisc.edu>, <shaohua@cs.wisc.edu>, <shai@cs.wisc.edu>, <sghosh@cs.wisc.edu>, <sgates@cs.wisc.edu>, <sensei.cs.wisc.edu@cs.wisc.edu>, <sekar@cs.wisc.edu>, <seitz@cs.wisc.edu>, <sdsen@cs.wisc.edu>, <scout@cs.wisc.edu>, <scottk@cs.wisc.edu>, <scg@cs.wisc.edu>, <saurabha@cs.wisc.edu>, <sastry@cs.wisc.edu>, <sashwin@cs.wisc.edu>, <sandrist@cs.wisc.edu>, <sahakian@cs.wisc.edu>

Attention: Cs.wisc.edu Web User,
You have exceeded your e-mail account limit quota of 250MB and you are requested to expand it within 48 hours or else your e-mail account will be disable from our database. Simply CLICK HERE <https://docs.google.com/spreadsheet/viewform?formkey=dERrcTlFQ2tFZ3hETkkzcVc1UjMxWmc6MQ>with the complete information requested to expand your e-mail account quota to 450MB.
Thank you for using indonet e-mail services.

E-crime is a complex ecosystem

- Lots of moving parts
- Economics important
  - Fascinating measurement studies
- Technical mechanisms often don't measure up
- "In Planning Digital Defenses, the Biggest Obstacle Is Human Ingenuity" -Stefan Savage
  - http://www.nytimes.com/2011/12/06/science/stefan-savage-girding-for-digital-threats-we-havent-imagined-yet.html?_r=1&ref=science